300-410 Implementing Cisco Enterprise Advanced Routing and Services (ENARSI)

Question #1 Topic 1

Refer to the exhibit. Users in the branch network of 2001:db8:0:4::/64 report that they cannot access the Internet.
Which command is issued in IPv6 router EIGRP 100 configuration mode to solve this issue?

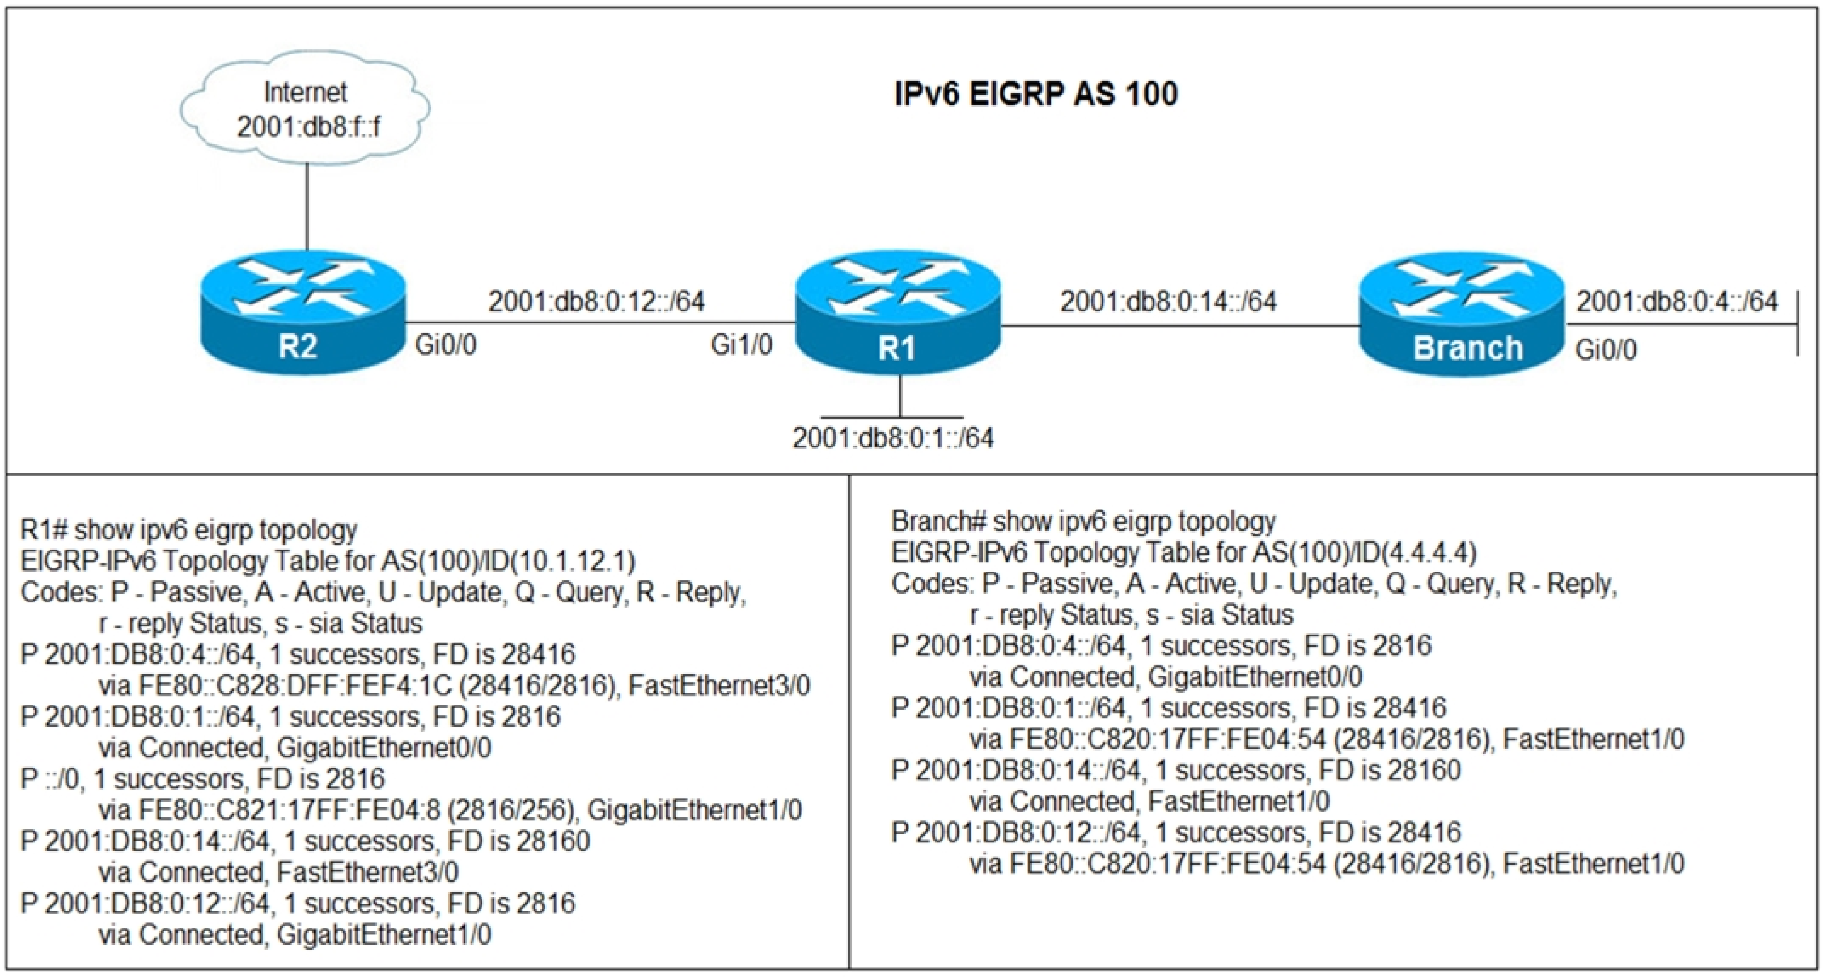

```
R1# show ipv6 eigrp topology
EIGRP-IPv6 Topology Table for AS(100)/ID(10.1.12.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status
P 2001:DB8:0:4::/64, 1 successors, FD is 28416
       via FE80::C828:DFF:FEF4:1C (28416/2816), FastEthernet3/0
P 2001:DB8:0:1::/64, 1 successors, FD is 2816
       via Connected, GigabitEthernet0/0
P ::/0, 1 successors, FD is 2816
       via FE80::C821:17FF:FE04:8 (2816/256), GigabitEthernet1/0
P 2001:DB8:0:14::/64, 1 successors, FD is 28160
       via Connected, FastEthernet3/0
P 2001:DB8:0:12::/64, 1 successors, FD is 2816
       via Connected, GigabitEthernet1/0
```

```
Branch# show ipv6 eigrp topology
EIGRP-IPv6 Topology Table for AS(100)/ID(4.4.4.4)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status
P 2001:DB8:0:4::/64, 1 successors, FD is 2816
       via Connected, GigabitEthernet0/0
P 2001:DB8:0:1::/64, 1 successors, FD is 28416
       via FE80::C820:17FF:FE04:54 (28416/2816), FastEthernet1/0
P 2001:DB8:0:14::/64, 1 successors, FD is 28160
       via Connected, FastEthernet1/0
P 2001:DB8:0:12::/64, 1 successors, FD is 28416
       via FE80::C820:17FF:FE04:54 (28416/2816), FastEthernet1/0
```

A. Issue the eigrp stub command on R1.

B. Issue the no eigrp stub command on R1.

C. Issue the eigrp stub command on R2.

D. Issue the no eigrp stub command on R2.

Correct Answer: *B*

Community vote distribution

B (100%)

dydzah Highly Voted 3 years, 5 months ago

B is correct. If you look closely you see that R1 learn the default route ::/0 from R2 via Gig0/1 but is not advertising it to Branch router

upvoted 17 times

khaganiabbasov Most Recent 2 months, 2 weeks ago

B correct

upvoted 1 times

Dacusai 7 months, 3 weeks ago

I lab it and it works just like the question says

upvoted 1 times

HungarianDish 7 months ago

I confirmed it with a lab, too. It's "B".

upvoted 1 times

MasterMatt 7 months, 3 weeks ago

Selected Answer: B

By observing the routing table we can determine that: 1) the default static route isn't learned on the branch router 2) R1 networks were learned. This concludes that on R1 we have eigrp stub connected which advertises only connected networks to Branch router.

upvoted 1 times

Wooker 9 months, 1 week ago

Selected Answer: B

If you look closely you see that R1 learns the default route::/0 from R2 via Gig0/1 but is not advertising it to the Branch router

Stub announces by default the connected and summary routes.

upvoted 1 times

baldebri 9 months, 2 weeks ago

https://1drv.ms/t/s!As7AtSYjWB003UPTilVcQVqEh-EC?e=LFktBj

upvoted 1 times

baldebri 9 months, 2 weeks ago

https://1drv.ms/i/s!As7AtSYjWB003UJBXO7GrdqdugRQ?e=o7Ea0B
the link to exhibt

upvoted 1 times

baldebri 9 months, 2 weeks ago

Refer to the diagram and the exhibit. All interfaces are participating in the routing processes shown in the diagram, and all neighborships have been formed. In addition, all the necessary routes have been exchanged. Which statement is correct in relationship to redistribution?
answers are :
BGP AS 65500 will not learn any EIGRP AS 100 prefixes.

Branch will learn all IPv4 prefixes except 192.0.2.1.

R2 will not learn any prefixes in EIGRP AS 100.

Branch will learn all IPv4 prefixes in the diagram.

upvoted 1 times

xziomal9 1 year, 8 months ago

The correct answer is: B

upvoted 2 times

Hack4 1 year, 10 months ago

Yes the given answer is correct

upvoted 1 times

Jenia1 1 year, 11 months ago

Selected Answer: B

The given answer is correct

upvoted 1 times

error_909 2 years, 2 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

ABELQF6 2 years, 5 months ago

Which is correct?

upvoted 1 times

Sadist1111 2 years, 10 months ago

Correct answer is B. In Router 1's RIB we can see the default route, which is advertised as an EIGRP Route. By-default Router can't advertised routes, which he leant vie EIGRP, when there is Stub configuration. So R1 can advertise gig1/0 link subnet, but not default route because of restriction.

upvoted 1 times

tomasz 2 years, 10 months ago

B is correct, stub eigrp router (in default) only send update about connected and summary routes. So R1 learned static route from R2 but didn't send this route to branch neighbor router

upvoted 1 times

CCNPWILL 2 years, 11 months ago

The answer is B.

The Branch rtr is missing a route due to not getting one. Stub means to not send a query downstream to find a particular IP addy. we need to disable that feature on R1. Answer is B.

upvoted 1 times

Question #2 Topic 1

Refer to the exhibit. Which configuration configures a policy on R1 to forward any traffic that is sourced from the 192.168.130.0/24 network to R2?

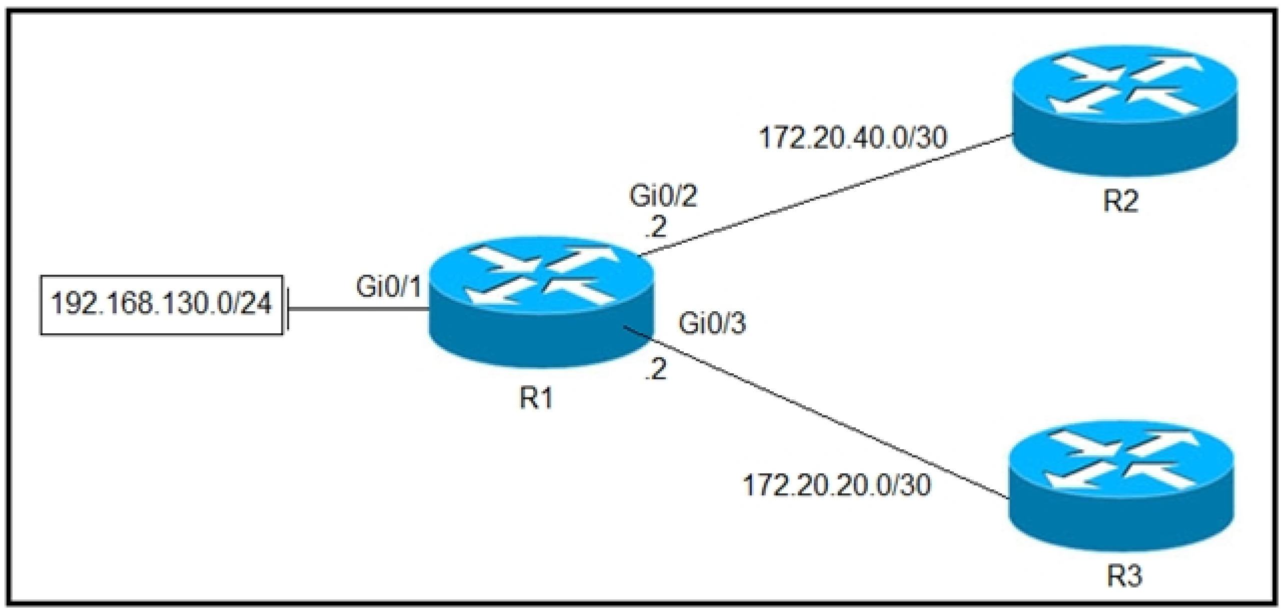

A.

```
access-list 1 permit 192.168.130.0 0.0.0.255
!
interface Gi0/2
ip policy route-map test
!
route-map test permit 10
match ip address 1
set ip next-hop 172.20.20.2
```

B.

```
access-list 1 permit 192.168.130.0 0.0.0.255
!
interface Gi0/1
ip policy route-map test
!
route-map test permit 10
match ip address 1
set ip next-hop 172.20.40.2
```

C.

```
access-list 1 permit 192.168.130.0 0.0.0.255
!
interface Gi0/2
ip policy route-map test
!
route-map test permit 10
match ip address 1
set ip next-hop 172.20.20.1
```

D.

```
access-list 1 permit 192.168.130.0 0.0.0.255
!
interface Gi0/1
ip policy route-map test
!
route-map test permit 10
match ip address 1
set ip next-hop 172.20.40.1
```

Correct Answer: *D*

⊟ **piiitrek** Highly Voted 2 years, 10 months ago

Answer is D - look at the address of the local router (R1) on p2p links - it has .2, so it means the next hop (the remote router) is .1

upvoted 9 times

⊟ **MasterMatt** Most Recent 7 months, 3 weeks ago

Common practice for an access list to be applied on an interface closest to the source. Also always set the next-hop on the adjacent IP (.1) on that subnet for the lookup.

upvoted 1 times

Alexloh 1 year, 6 months ago

Answer D is correct.

upvoted 1 times

xziomal9 1 year, 8 months ago

The correct answer is: D

upvoted 1 times

andrew230 2 years, 2 months ago

D is correct

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Wesgo 2 years, 8 months ago

To clear out the confusion if B or D, it is D indeed. See Step 5 below: "Specifies the action to be taken on the packets that match the criteria. Sets next hop to which to route the packet (the next hop must be adjacent)."

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960x/software/15_2_6_e/configuration_guide/b_1526e_consolidated_2960x_cg/b_1526e_consolidated_2960x_cg_chapter_01010000.pdf

upvoted 1 times

uglyprawn 2 years, 9 months ago

correct answer is D. a similar example is in the enarsi book page 624

upvoted 1 times

Pys17 2 years, 9 months ago

Answer D in correct.

upvoted 3 times

Benzzyy 2 years, 10 months ago

Answer is actually B.
The Next-Hop IP Address should be the upstream routers IP address for that link.

upvoted 1 times

> **AliMo123** 2 years, 4 months ago
>
> 172.20.40.1 is the next hop IP for R1 not 172.20.40.2 which is one of the R1 interface itsels.
>
> upvoted 2 times
>
> **Pb1805** 2 years, 7 months ago
>
> Yes. Thats why correct answer is D
>
> upvoted 1 times

tomasz 2 years, 10 months ago

B obviously...

upvoted 1 times

Question #3 Topic 1

R2 has a locally originated prefix 192.168.130.0/24 and has these configurations:

```
ip prefix-list test seq 5 permit 192.168.130.0/24
!
route-map OUT permit10
match ip address prefix-list test
set as-path prepend 65000
```

What is the result when the route-map OUT command is applied toward an eBGP neighbor R1 (1.1.1.1) by using the neighbor 1.1.1.1 route-map OUT out command?

A. R1 sees 192.168.130.0/24 as two AS hops away instead of one AS hop away.

B. R1 does not accept any routes other than 192.168.130.0/24

C. R1 does not forward traffic that is destined for 192.168.30.0/24

D. Network 192.168.130.0/24 is not allowed in the R1 table

Correct Answer: *A*

Community vote distribution

A (78%) B (22%)

⊟ **Guitarman** Highly Voted 3 years, 4 months ago

I'm going with A. The as-prepend will add the additional AS identifier which in turn makes the route 2 AS hops a way. This is used with multihomed ISP configurations to determine the path of incoming traffic.

upvoted 12 times

⊟ **HungarianDish** Most Recent 8 months ago

Selected Answer: A

https://community.cisco.com/t5/networking-knowledge-base/understanding-bgp-best-path-selection-manipulation/ta-p/3150576

upvoted 3 times

⊟ **rogabor81** 11 months ago

Selected Answer: B

I would say B. Who said that the ebgp peers are directly connected? it can be an ebgp-multihop 3 or something in the config. The only answer what is right in any circumstances is B....

upvoted 2 times

⊟ **Almylle** 6 months, 3 weeks ago

is an "OUT" route map, so u are advertising only the 192.168.130.0/24, so it cannot be the Answer B.

upvoted 2 times

⊟ **nicoaburto** 11 months, 2 weeks ago

Selected Answer: A

A - PREPEND 65000 in the as-path, R2 see 65000 65000 for this prefix

upvoted 1 times

⊟ **MD_Shox** 1 year ago

A. R1 sees 192.168.130.0/24 as two AS hops away instead of one AS hop away.
and R2 does filter all other route adverticements other than 192.168.130.0/24 when sending to R1, fue to ipmlicit deny (missing route-map permit 20 statement

upvoted 1 times

⊟ **kaisehhop** 1 year, 1 month ago

The given answer is correct

upvoted 1 times

⊟ **bryaberson** 1 year, 2 months ago

What if the Routemap does not have a permit statement sequence 20? Then B should also be an answer as the explicit deny statement will deny any network other than 192.168.130.x

upvoted 2 times

⊟ **potato_inet0** 7 months ago

The wording is tricky here, R1 will accept routes other than 192.168.130.x because R1 does not have any RM in place, R2 however will not sent any routes other than 192.168.130.x

upvoted 2 times

⊟ **Alexloh** 1 year, 5 months ago

```
//ORIGINAL WITHOUT AS-PREPEND//
R3#sh ip bgp | i 192.
*> 192.168.130.0 2.2.2.2 0 65002 65000 i

//ORIGINAL WITH AS-PREPEND//
R3#sh ip bgp | i 192.
*> 192.168.130.0 2.2.2.2 0 65002 65000 65000 i
```

upvoted 1 times

Alexloh 1 year, 6 months ago

Selected Answer: A

A is correct

upvoted 2 times

xziomal9 1 year, 8 months ago

The correct answer is: A

upvoted 1 times

Hack4 1 year, 10 months ago

the given answer is correct

upvoted 1 times

Networkingguy 1 year, 11 months ago

Selected Answer: A

A looks to be correct

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

uglyprawn 2 years, 9 months ago

i would go with a as well. we prepend as to make an as path longer to influence incoming traffic to take the other path(shorter as_path)

upvoted 1 times

Os_ 3 years, 4 months ago

What is the result when the route-map OUT command is applied toward an eBGP neighbor R1 (1.1.1.1) by using the neighbor 1.1.1.1 route-map OUT out command? What??

upvoted 1 times

> **thissiteisgreat** 2 years, 11 months ago
>
> Because the commands are issued on R2, and its neighbor, R1 is receiving the appended route. So, to R1, the AS Path attribute for the route is 2 hop away
>
> upvoted 4 times

Question #4 Topic 1

Which method changes the forwarding decision that a router makes without first changing the routing table or influencing the IP data plane?

A. nonbroadcast multiaccess

B. packet switching

C. policy-based routing

D. forwarding information base

Correct Answer: *C*

Community vote distribution

C (89%) 11%

T_Cos 1 month ago

C is correct

upvoted 1 times

siscoFe 5 months, 4 weeks ago

PBR takes precedence from Routing table when it comes to routing decisions iff it is configured already. So it makes sense that it is answered as C.

upvoted 1 times

Wooker 9 months, 1 week ago

Selected Answer: C

PBR is the method to influence route without changin any on RIB.

upvoted 1 times

Koume 11 months, 1 week ago

Selected Answer: C

PBR is the method to influence route without changin any on RIB.

upvoted 1 times

Noproblem22 1 year, 1 month ago

Selected Answer: C

B is the best answer

upvoted 1 times

DumpsterFire 1 year, 3 months ago

Selected Answer: C

C is correct

upvoted 2 times

Router 1 year, 3 months ago

c is the correct ans

upvoted 1 times

Alexloh 1 year, 6 months ago

Selected Answer: C

C is correct.

upvoted 1 times

xziomal9 1 year, 8 months ago

The correct answer is: C

upvoted 1 times

Nhan 1 year, 9 months ago

The given answer correct

upvoted 1 times

Baiji 1 year, 10 months ago

Selected Answer: C

Answer looks to be C here

upvoted 1 times

Networkingguy 1 year, 11 months ago

Selected Answer: C

Answer looks to be C here

upvoted 1 times

[Removed] 1 year, 10 months ago

Its D. The key word in this question is first lol... Without FIRST, meaning its changing whatever follows. After the routing table changes the FIB will update and the router will use that for the forwarding decision. Cisco giving us a english exam as well smh..

upvoted 1 times

[Removed] 1 year, 11 months ago

Selected Answer: D

So the only answer here that changes the routing table is D, and the router does use that to make forwarding decisions after all so im going with D.

upvoted 1 times

thegolden3 1 year, 11 months ago

changes the forwarding decision that a router makes WITHOUTfirst changing the routing table, answer is C.

upvoted 3 times

[Removed] 1 year, 10 months ago

Lol the way its worded can be interpreted 2 different ways.

upvoted 1 times

[Removed] 1 year, 10 months ago

This is one of those English type questions. You really have to comprehend what they're asking. Without first, meaning changing these things first. That points to the FIB. Any changes made to the routing table/data plane populates the FIB which the routers uses to forward. PBR is thrown in there for confusion because the first portion of the question it fits but then it doesn't make sense once you read the entire question.

upvoted 1 times

bogd 1 year, 10 months ago

No, "without first changing x" means that it DOES NOT change X. Yes, it is "one of those English type questions", but you are misinterpreting it here...

upvoted 2 times

Koume 11 months, 1 week ago

I think is a spooky question but, i really go for C, as if you take all the sentece says "without first changing the routing table or influencing the IP data plane".
when talks the routing table speak as the RIB, but also the Dataplane is clearly talking about the FIB. I understood that means without ifluencing both. and the only one that do tha is PBR.

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

uglyprawn 2 years, 9 months ago

answer is c. pbr does not affect make changes in rib

upvoted 1 times

chub 3 years, 2 months ago

PBR doesn't change the routing table.

upvoted 4 times

CCNPWILL 2 years, 11 months ago

Agreed. HIGHLY recommend they change the wording here.

upvoted 1 times

CCNPWILL 2 years, 11 months ago

The answer is still C though.

upvoted 1 times

Question #5 *Topic 1*

Refer to the exhibits. The output of the trace route from R5 shows a loop in the network.

Which configuration prevents this loop?

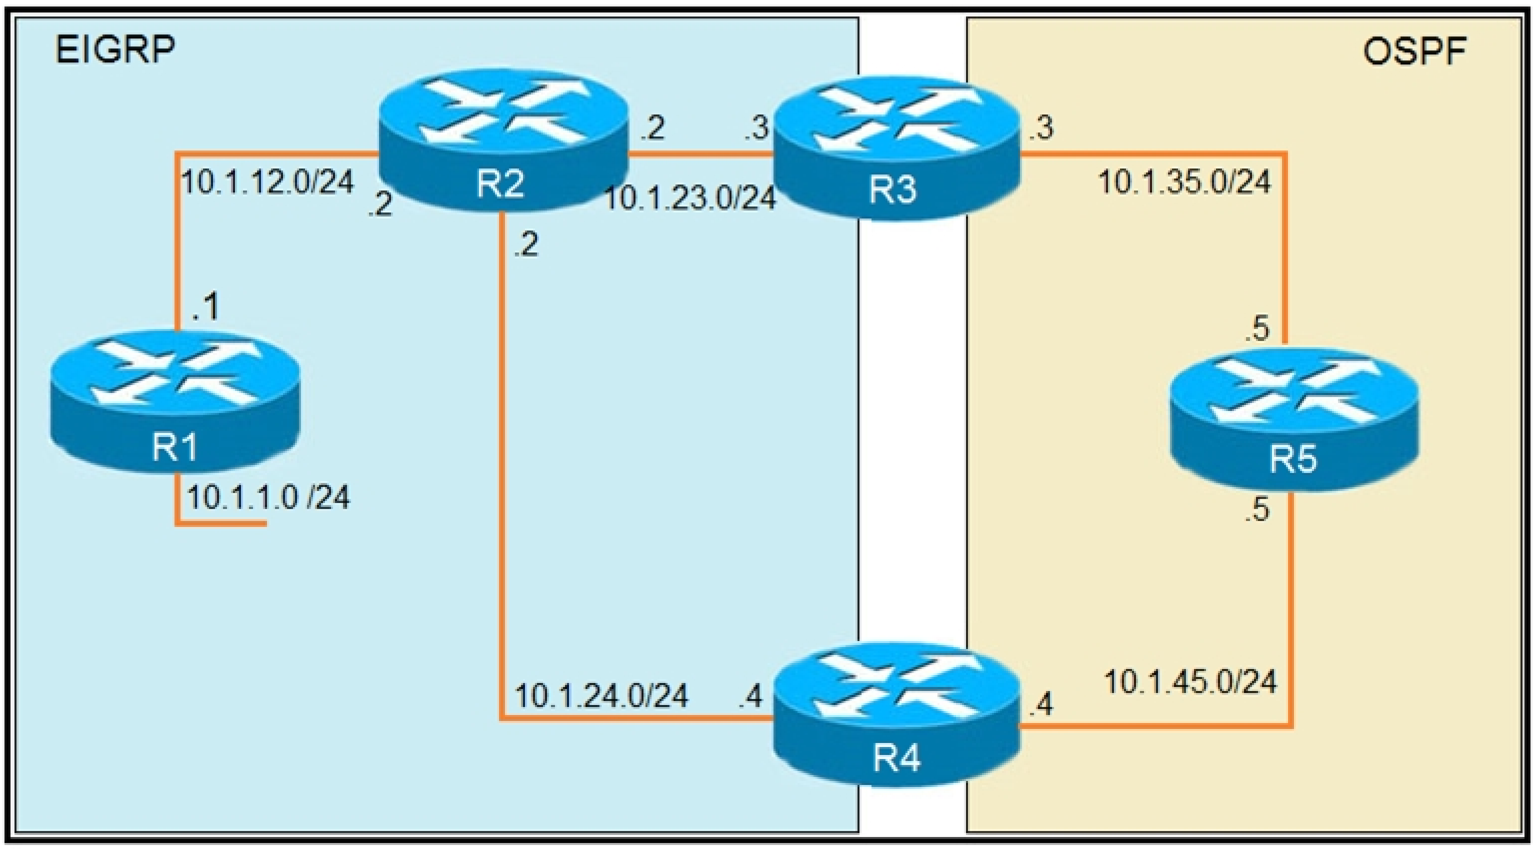

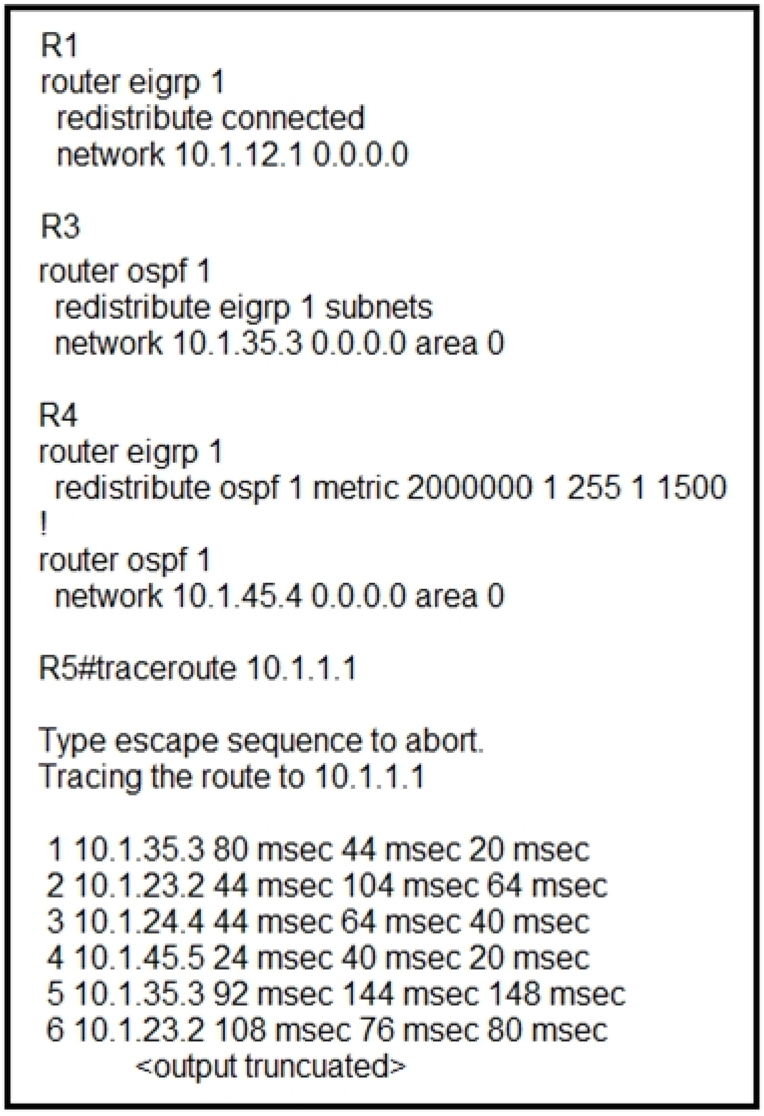

```
R1
router eigrp 1
 redistribute connected
 network 10.1.12.1 0.0.0.0

R3
router ospf 1
 redistribute eigrp 1 subnets
 network 10.1.35.3 0.0.0.0 area 0

R4
router eigrp 1
 redistribute ospf 1 metric 2000000 1 255 1 1500
!
router ospf 1
 network 10.1.45.4 0.0.0.0 area 0

R5#traceroute 10.1.1.1

Type escape sequence to abort.
Tracing the route to 10.1.1.1

 1 10.1.35.3 80 msec 44 msec 20 msec
 2 10.1.23.2 44 msec 104 msec 64 msec
 3 10.1.24.4 44 msec 64 msec 40 msec
 4 10.1.45.5 24 msec 40 msec 20 msec
 5 10.1.35.3 92 msec 144 msec 148 msec
 6 10.1.23.2 108 msec 76 msec 80 msec
        <output truncated>
```

A.

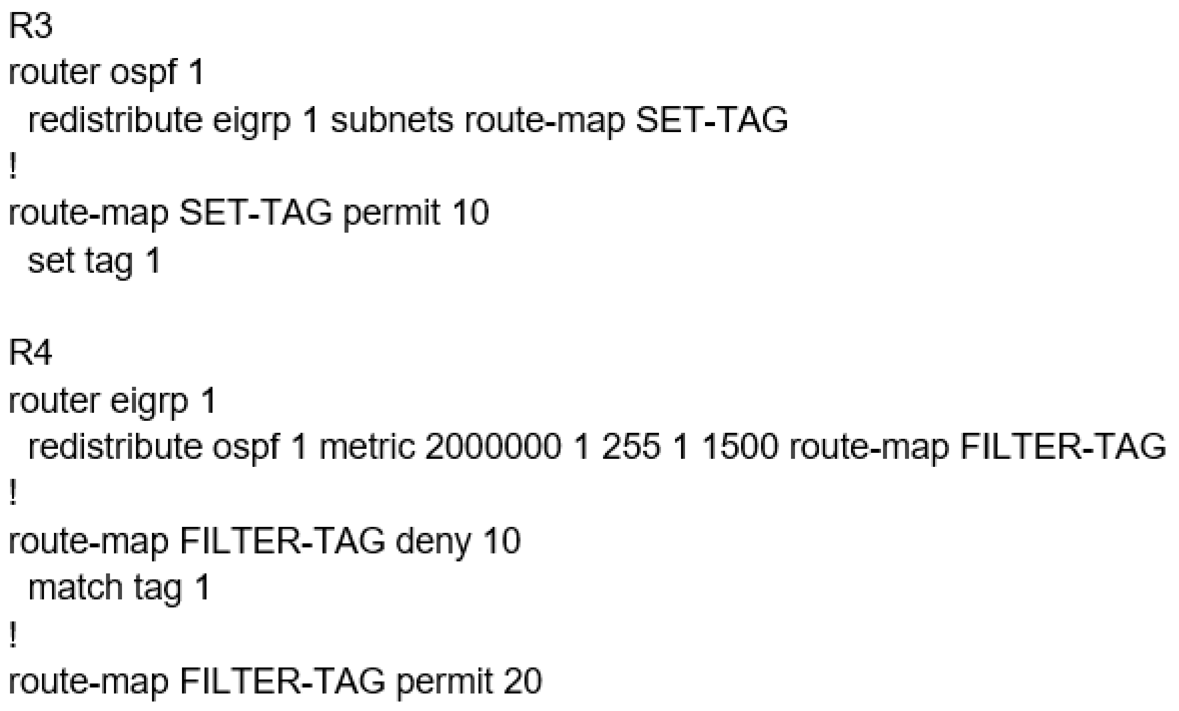

```
R3
router ospf 1
 redistribute eigrp 1 subnets route-map SET-TAG
!
route-map SET-TAG permit 10
 set tag 1

R4
router eigrp 1
 redistribute ospf 1 metric 2000000 1 255 1 1500 route-map FILTER-TAG
!
route-map FILTER-TAG deny 10
 match tag 1
!
route-map FILTER-TAG permit 20
```

B.

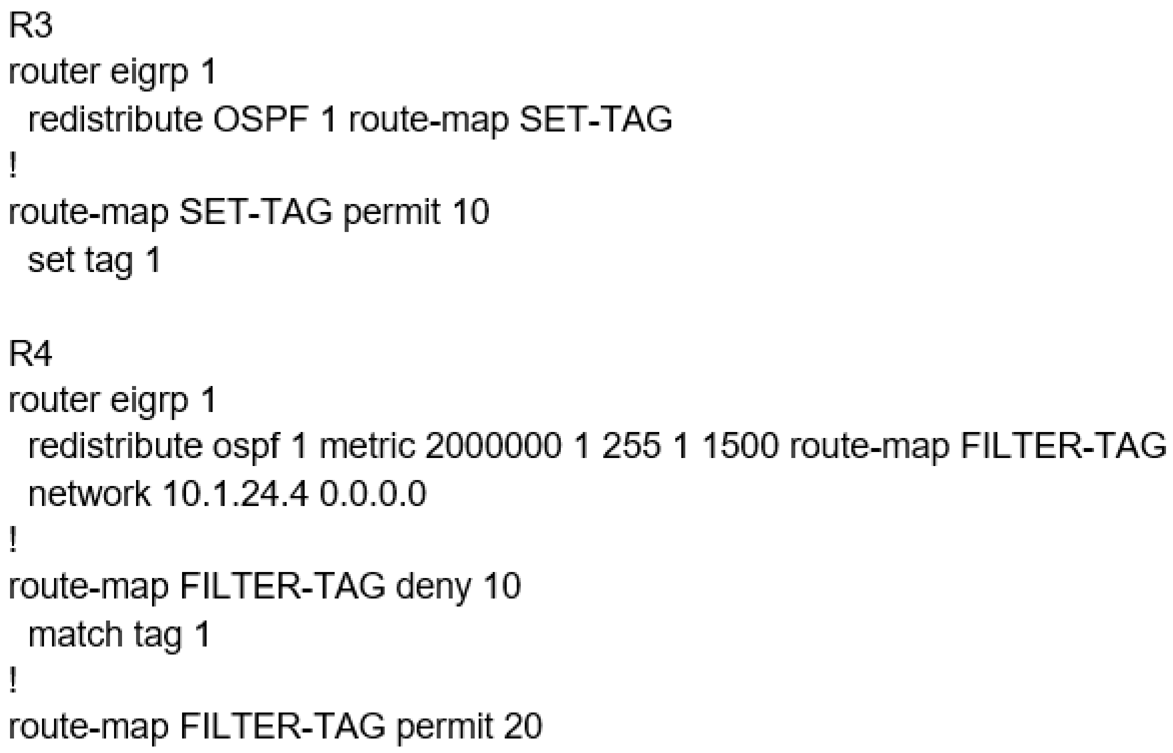

```
R3
router eigrp 1
  redistribute OSPF 1 route-map SET-TAG
!
route-map SET-TAG permit 10
  set tag 1

R4
router eigrp 1
  redistribute ospf 1 metric 2000000 1 255 1 1500 route-map FILTER-TAG
  network 10.1.24.4 0.0.0.0
!
route-map FILTER-TAG deny 10
  match tag 1
!
route-map FILTER-TAG permit 20
```

C.

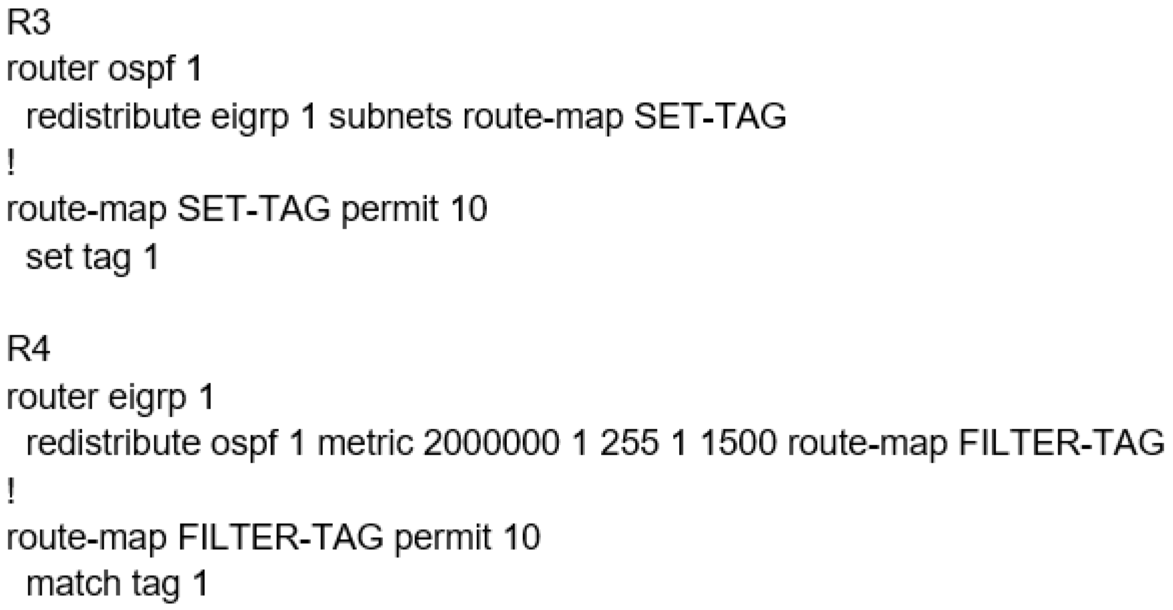

```
R3
router ospf 1
  redistribute eigrp 1 subnets route-map SET-TAG
!
route-map SET-TAG permit 10
  set tag 1

R4
router eigrp 1
  redistribute ospf 1 metric 2000000 1 255 1 1500 route-map FILTER-TAG
!
route-map FILTER-TAG permit 10
  match tag 1
```

D.

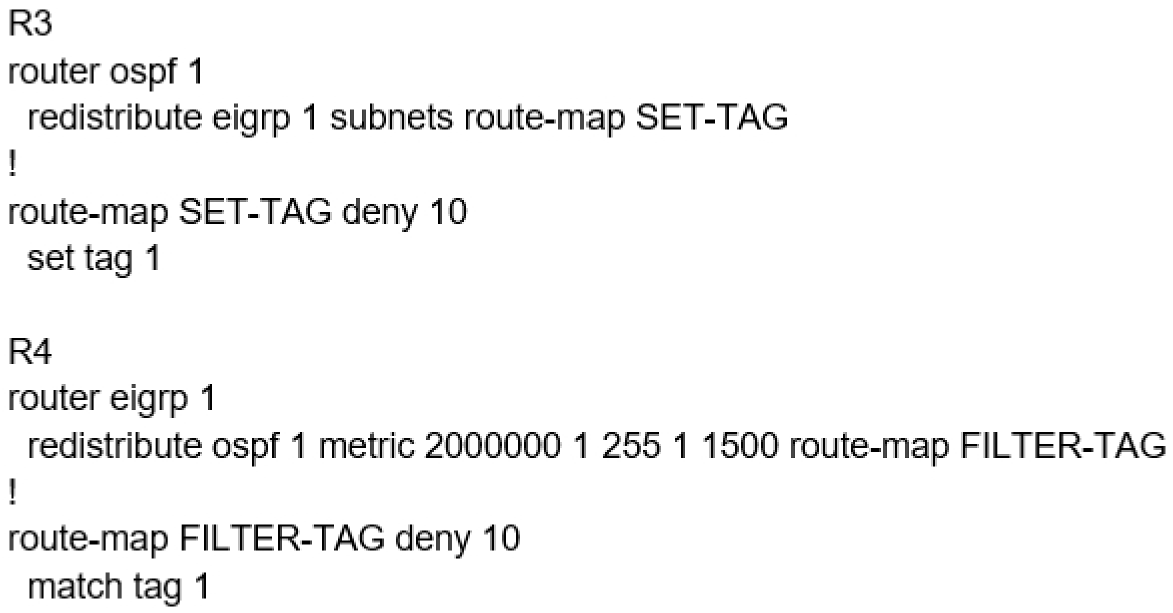

```
R3
router ospf 1
  redistribute eigrp 1 subnets route-map SET-TAG
!
route-map SET-TAG deny 10
  set tag 1

R4
router eigrp 1
  redistribute ospf 1 metric 2000000 1 255 1 1500 route-map FILTER-TAG
!
route-map FILTER-TAG deny 10
  match tag 1
```

Correct Answer: *B*

⊟ **Deadliftn** Highly Voted 1 year, 4 months ago

Answer is A but the available answers are all written wrong either way. Whoever wrote this is crazy. But, the CLOSEST possible answer would be A. Whoever writes questions for the Cisco exams are absolutely ignorant in how they write questions OR they are being deliberate in trying to fool the test takers, which is sad.

upvoted 12 times

⊟ **Koume** 11 months, 1 week ago

After a thoght analysis is not B, Il explain Why, R3 are redistributing OSPF into EIGRP and setting the tag 1, but notice that the tag 1 is being announced on EIGRP process, when R4 redistribute OSPF into EIGRP with the route map it will not match anything because that tag is no been announce by ospf process.
So on R4 the R1 network will be redistributed back and being announced to R2, as the reported distance reset by redistribution then when packet arrives to R2 the R4 router will be prefered.
In conclusion B is not Correct, the most closest is A

upvoted 4 times

⊟ **net_eng10021** 2 months, 2 weeks ago

I see the same thing as Koume has described above. The network is not tagged in the ospf domain.

upvoted 1 times

HungarianDish Highly Voted 6 months, 3 weeks ago

"A"

I have redone this lab. Introduced the loop, then applied solution "A". It did actually prevented the loop.

Before applying "A":

```
R5#trac 10.1.1.1
Type escape sequence to abort.
Tracing the route to 10.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
1 10.1.35.3 2 msec 1 msec 2 msec
2 10.1.23.2 2 msec 2 msec 2 msec
3 10.1.24.4 2 msec 2 msec 2 msec
4 10.1.45.5 1 msec 2 msec 2 msec
5 10.1.35.3 3 msec 2 msec 2 msec
6 10.1.23.2 3 msec 2 msec 3 msec
7 10.1.24.4 3 msec 3 msec 3 msec
8 10.1.45.5 2 msec 2 msec 2 msec
9 10.1.35.3 4 msec 3 msec 3 msec
10 10.1.23.2 3 msec 4 msec 4 msec
<output ommited>
```

After applying "A":

```
R5#trac 10.1.1.1
Type escape sequence to abort.
Tracing the route to 10.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
1 10.1.35.3 2 msec 2 msec 1 msec
2 10.1.23.2 2 msec 2 msec 2 msec
3 10.1.12.1 2 msec * 2 msec
R5#
```

upvoted 6 times

HungarianDish 6 months, 3 weeks ago

Before applying "A" - 10.1.1.0/24 is learned from OSPF:

```
R4#sh ip eigrp 1 top 10.1.1.0/24 | sec External
Composite metric is (2816/0), route is External
External data:
AS number of route is 1
External protocol is OSPF, external metric is 20
Administrator tag is 1 (0x00000001)
```

After applying "A" - tagged ospf routes are filtered, 10.1.1.0/24 is learned from redistribute connected via eigrp:

```
R4#sh ip eigrp 1 top 10.1.1.0/24 | sec External
Composite metric is (131072/130816), route is External
External data:
AS number of route is 0
External protocol is Connected, external metric is 0
Administrator tag is 0 (0x00000000)
```

upvoted 1 times

HungarianDish 6 months, 1 week ago

Before applying solution A, R2 sees two redistributed routes in eigrp, one from redistribute connected, and another from redistribute ospf. R2 trusts ospf more, and sends traffic to R4. Loop is created.

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Prefix is tagged:

```
R4#sh ip route 10.1.1.1
Routing entry for 10.1.1.0/24
Known via "ospf 1", distance 110, metric 20
Tag 1, type extern 2, forward metric 2
Redistributing via eigrp 1

R4#sh run | sec router eigrp
router eigrp 1
network 10.1.24.0 0.0.0.255
redistribute ospf 1 metric 1000000 1 255 1 1500 route-map FILTER-TAG
R4#
R4#sh run | sec route-map
redistribute ospf 1 metric 1000000 1 255 1 1500 route-map FILTER-TAG
route-map FILTER-TAG deny 10
match tag 1
route-map FILTER-TAG permit 20
```

upvoted 1 times

net_eng10021 Most Recent 2 months, 2 weeks ago

I like A here. The problem with B is that the 10.1.1.0/24 subnet is not getting tagged on the eigrp to ospf redistribution at R3. Hence, R4, can't block it from on the ospf to eigrp redistribution at R4.

upvoted 1 times

Mohammad963 3 months, 2 weeks ago

I'll go with A, 100% .

upvoted 2 times

LanreDipeolu 3 months, 2 weeks ago

B is the correct answer from the fact that R4 advertised the important route of 10.1.24.4, which other options did not. Also technically set tag1 in R3 and denied it in R4.

upvoted 1 times

Chiaretta 4 months, 3 weeks ago

The right answer is A.

upvoted 2 times

inteldarvid 5 months ago

100%% option "A"

upvoted 1 times

HungarianDish 7 months ago

For me also "A" seems to be the closest, because it is applying the tag on the correct combination of protocol & router. I labbed this scenario in CML, but I was unable to reproduce a loop with this configuration.

upvoted 1 times

AinsB 7 months, 1 week ago

Answer is B. R1 is advertising the connected 10.1.1.0 as an external network AD 170. OSPF advertises it as 110 so by default R4 will take the path through R5->R3 to get to 10.1.1.0. R2 IS advertising it at 170 to R4 even though it is a shorter path. If we block advertisement from R5 for this network then the better path of R4 -> R2 will be chosen.

upvoted 1 times

Dacusai 8 months ago

In answer C and D is missing the permit 20 on the route map mining that no other routes will be added to the routing table and one of them has a permit so it still has the loop.

upvoted 1 times

Dacusai 8 months ago

According the configuration in R3 you redistribute EIGRP into OSPF and answer B say other wise, so A is the correct one.

upvoted 1 times

anonymous1966 8 months, 3 weeks ago

***(A) is correct
Blue ---> Yelow (SET a tag)
Blue <--- Yelow (BLOCK tagged updates back and permit not tagged)

(B) wrong direction back
Yelow --> Blue (SET a tag)
Yelow --> Blue (BLOCK tagged updates and permit not tagged)

(C) a permit statement is missing
Blue ---> Yelow (SET a tag)
Blue <--- Yelow (BLOCK tagged updates back and DO NOT permit not tagged)

(D) a deny statement is used wrongly
Blue ---> Yelow (DO NOT SET a tag)
Blue <--- Yelow (BLOCK tagged updates back and DO NOT permit not tagged)

upvoted 4 times

> **[Removed]** 4 months ago
>
> beautiful explanation.
>
> upvoted 1 times

Wooker 9 months, 1 week ago

correct answer is A.

upvoted 1 times

M_Abdulkarim 1 year, 4 months ago

Correct Answer is A

upvoted 3 times

Edwinmolinab 1 year, 5 months ago

I tested the options in GNS3 and the only one that works was option B. Given answer is correct

upvoted 1 times

> **Edwinmolinab** 1 year, 5 months ago
>
> I was wrong the correct answer is A the tag go in the ospf direction on R3 because R4 waits tagging 1 to prevent the loop
>
> upvoted 4 times

Reikidude00 1 year, 5 months ago

A is de correct answer.
In B there is no redist into OSPF
upvoted 2 times

timtgh 1 year, 7 months ago

How can A be correct if EIGRP doesn't have a network command on R4?
upvoted 2 times

Dominik_Networker 10 months ago

It has, but it is omitted from the picture. Cisco exams questions are unfortunately sometimes written in a way, that you need to complete the missing parts of the question too, by looking at details. In this case it doesn't show the network command, but it shows, that it is in fact there by using the traceroute. If it wouldn't be there, than the network on the exhibit wouldn't have a loop, but it would be a one way only communication between the OSPF and EIGRP. I hope this isn't too confusing, but my conclusion is, that the command is there, just not shown on the picture, just to confuse the test taker
upvoted 2 times

Question #6 *Topic 1*

Refer to the exhibit. An engineer configures a static route on a router, but when the engineer checks the route to the destination, a different next hop is chosen.
What is the reason for this?

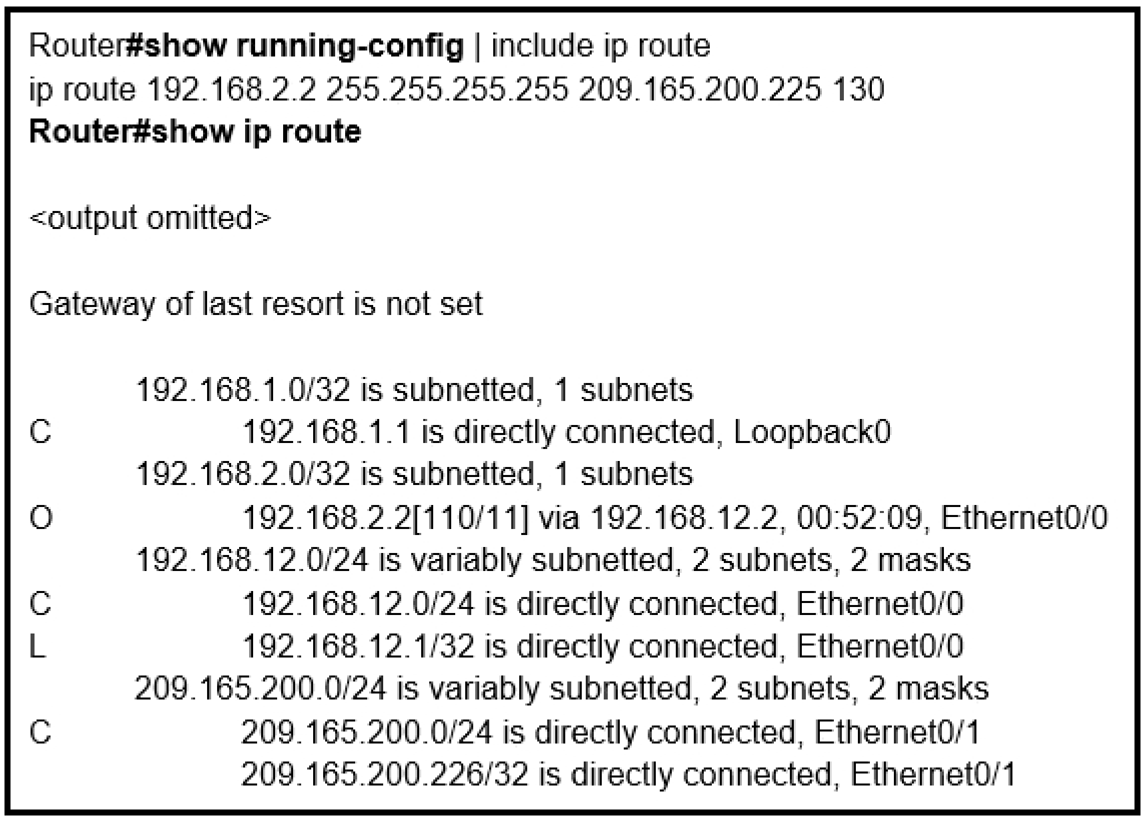

```
Router#show running-config | include ip route
ip route 192.168.2.2 255.255.255.255 209.165.200.225 130
Router#show ip route

<output omitted>

Gateway of last resort is not set

        192.168.1.0/32 is subnetted, 1 subnets
C               192.168.1.1 is directly connected, Loopback0
        192.168.2.0/32 is subnetted, 1 subnets
O               192.168.2.2[110/11] via 192.168.12.2, 00:52:09, Ethernet0/0
        192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks
C               192.168.12.0/24 is directly connected, Ethernet0/0
L               192.168.12.1/32 is directly connected, Ethernet0/0
        209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
C               209.165.200.0/24 is directly connected, Ethernet0/1
                209.165.200.226/32 is directly connected, Ethernet0/1
```

A. Dynamic routing protocols always have priority over static routes.

B. The metric of the OSPF route is lower than the metric of the static route.

C. The configured AD for the static route is higher than the AD of OSPF.

D. The syntax of the static route is not valid, so the route is not considered.

Correct Answer: *C*

Community vote distribution

C (100%)

Malasxd 7 months, 2 weeks ago

Selected Answer: C

The correct answer is: C

upvoted 1 times

Koume 11 months, 1 week ago

Selected Answer: C

Clearly is C, as AD of static routes is 130 vs 110 on ospf

upvoted 1 times

Alexloh 1 year, 6 months ago

Selected Answer: C

C is correct because the AD for the static route was set to 130 vs. OSPD default 110.

upvoted 2 times

xziomal9 1 year, 8 months ago

The correct answer is: C

upvoted 1 times

jester_2020 1 year, 8 months ago

C is correct. The AD of static route by default is lower (0) than OSPF (110) but the example shows it was override to 130.

upvoted 1 times

Networkingguy 1 year, 11 months ago

Selected Answer: C

C is correct here

upvoted 1 times

Nonono 1 year, 11 months ago

Selected Answer: C

Answer is correct

upvoted 1 times

andrew230 2 years, 2 months ago

C is correct ,the AD for static route is 130 ,the AD in OSPF is 110; 130 > 110 so win OSPF

upvoted 1 times

error_909 2 years, 3 months ago

The configured AD for the static route is higher than the AD of OSPF.

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Wesgo 2 years, 8 months ago

This would be easy for CCNA too

upvoted 1 times

Benzzyy 2 years, 10 months ago

C is correct

upvoted 2 times

Question #7

Topic 1

Refer to the exhibit. An engineer is trying to generate a summary route in OSPF for network 10.0.0.0/8, but the summary route does not show up in the routing table.

Why is the summary route missing?

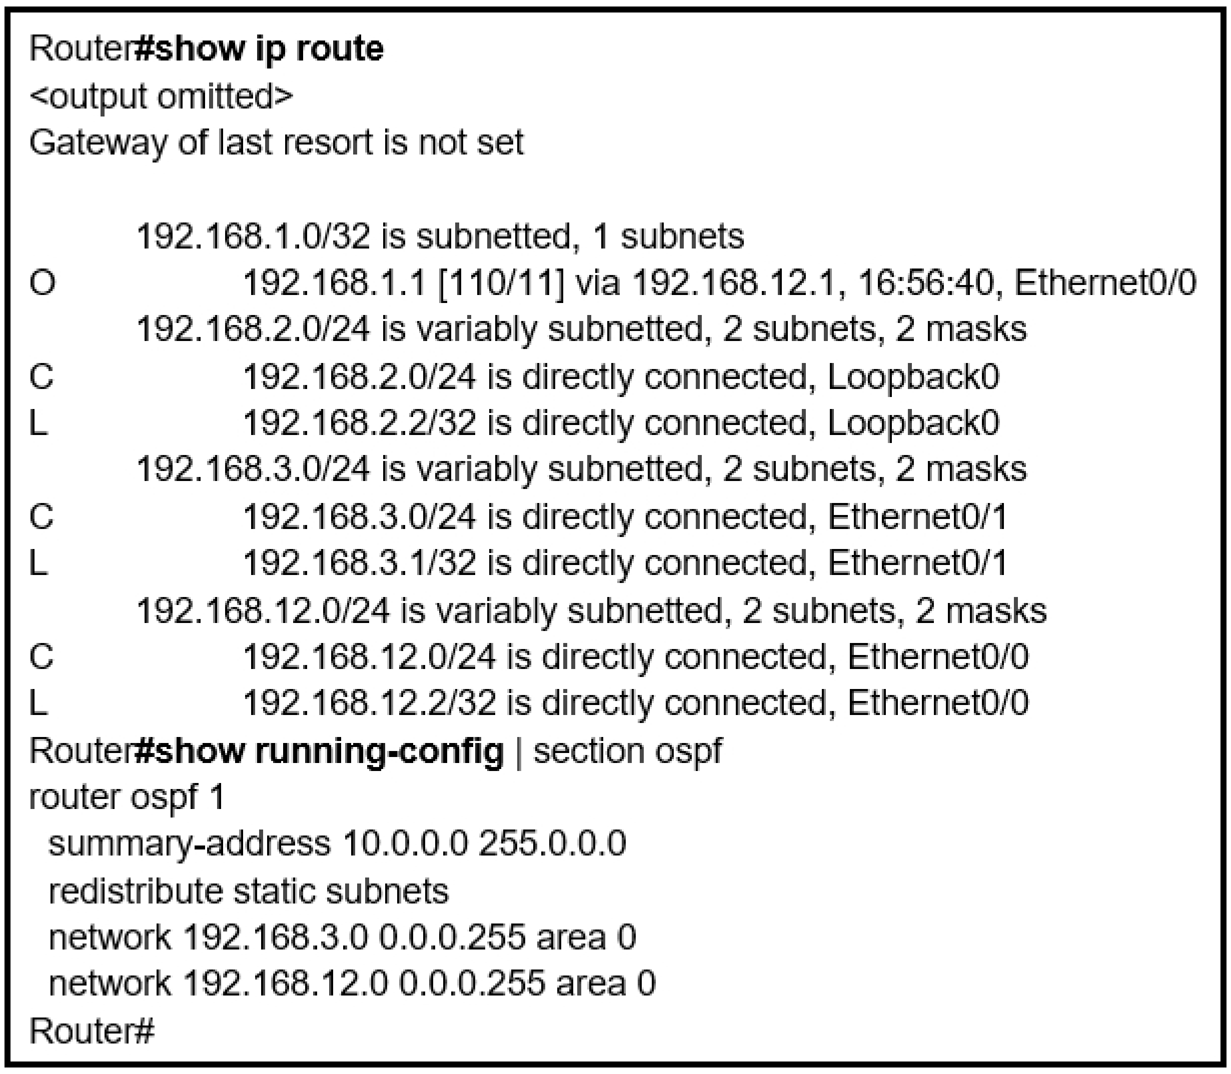

```
Router#show ip route
<output omitted>
Gateway of last resort is not set

      192.168.1.0/32 is subnetted, 1 subnets
O           192.168.1.1 [110/11] via 192.168.12.1, 16:56:40, Ethernet0/0
      192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
C           192.168.2.0/24 is directly connected, Loopback0
L           192.168.2.2/32 is directly connected, Loopback0
      192.168.3.0/24 is variably subnetted, 2 subnets, 2 masks
C           192.168.3.0/24 is directly connected, Ethernet0/1
L           192.168.3.1/32 is directly connected, Ethernet0/1
      192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks
C           192.168.12.0/24 is directly connected, Ethernet0/0
L           192.168.12.2/32 is directly connected, Ethernet0/0
Router#show running-config | section ospf
router ospf 1
 summary-address 10.0.0.0 255.0.0.0
 redistribute static subnets
 network 192.168.3.0 0.0.0.255 area 0
 network 192.168.12.0 0.0.0.255 area 0
Router#
```

A. The summary-address command is used only for summarizing prefixes between areas.

B. The summary route is visible only in the OSPF database, not in the routing table.

C. There is no route for a subnet inside 10.0.0.0/8, so the summary route is not generated.

D. The summary route is not visible on this router, but it is visible on other OSPF routers in the same area.

Correct Answer: *C*

Community vote distribution

C (100%)

conft 4 months ago

given answer is the correct.

upvoted 1 times

Malasxd 7 months, 2 weeks ago

Selected Answer: C

the command "summary-address" is used to summary external routes (O E1/2) in ASBR. The command "redistribute static" in the question makes the router a ASBR.

To summary inter area routes into ABRs you use "area x range" command.

In both cases the summary route is advertised only if the RIB has a route that matches the summary prefix.

upvoted 3 times

Alexloh 1 year, 5 months ago

I have tested in lab, the summary-address only worked if you have the valid route on your routing table.

upvoted 3 times

tefacert 1 year, 7 months ago

What about A? this is not an ABR, it only has area 0

upvoted 1 times

timtgh 1 year, 7 months ago

This summary command is not for ABRs, it's for ASBRs, so Option A is wrong.

upvoted 2 times

[Removed] 1 year, 10 months ago

The given answer is correct

upvoted 1 times

Jenia1 1 year, 11 months ago

Selected Answer: C

The given answer is correct

upvoted 1 times

Stivostine 2 years ago

Selected Answer: C

C is ok

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

Question #8 *Topic 1*

Refer to the exhibit. An engineer is trying to block the route to 192.168.2.2 from the routing table by using the configuration that is shown. The route is still present in the routing table as an OSPF route.

Which action blocks the route?

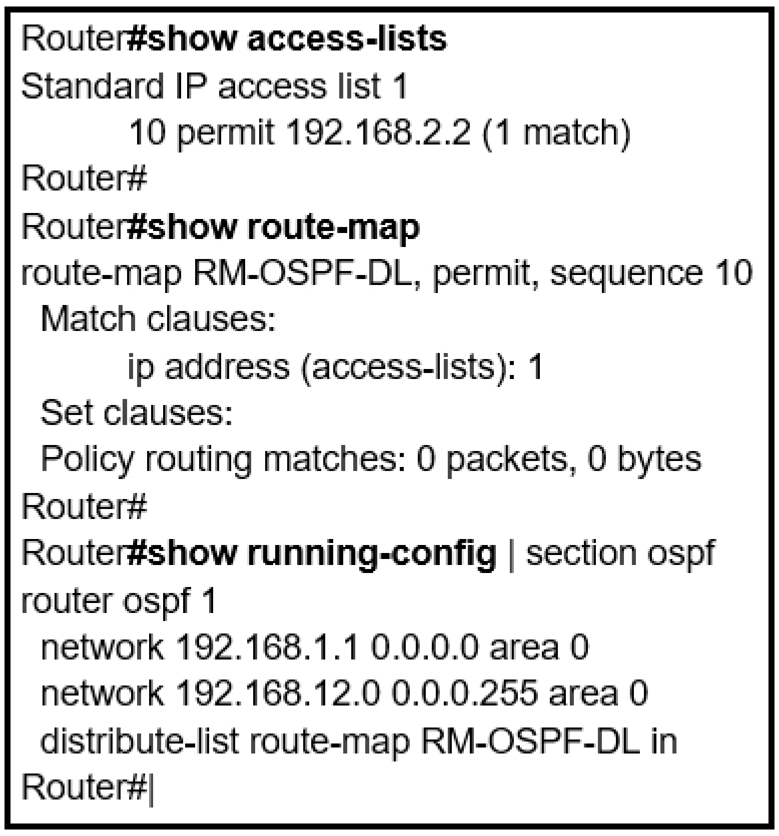

```
Router#show access-lists
Standard IP access list 1
        10 permit 192.168.2.2 (1 match)
Router#
Router#show route-map
route-map RM-OSPF-DL, permit, sequence 10
  Match clauses:
        ip address (access-lists): 1
  Set clauses:
  Policy routing matches: 0 packets, 0 bytes
Router#
Router#show running-config | section ospf
router ospf 1
  network 192.168.1.1 0.0.0.0 area 0
  network 192.168.12.0 0.0.0.255 area 0
  distribute-list route-map RM-OSPF-DL in
Router#|
```

A. Use an extended access list instead of a standard access list.

B. Change sequence 10 in the route-map command from permit to deny.

C. Use a prefix list instead of an access list in the route map.

D. Add this statement to the route map: route-map RM-OSPF-DL deny 20.

Correct Answer: *C*

Community vote distribution

B (81%) C (19%)

TigerDrev Highly Voted 3 years, 6 months ago

Agree with B

upvoted 16 times

ALONZINGER Highly Voted 3 years, 6 months ago

should be B i think

upvoted 13 times

LI123123 Most Recent 2 months ago

Selected Answer: B

I choose B

upvoted 1 times

jansan55 3 months, 2 weeks ago

Selected Answer: B

Tested in lab.
Answer A: permit in ACL and permit in route-map - 192.168.2.2 remain in the routing table.
Answer B: deny in ACL and permit in route-map will remove 192.168.2.2 from the routing table.
Answer C: permit in prefix-list and permit in route-map - 192.168.2.2 remain in the routing table.
Answer D: the sequence 10 already let the 192.168.2.2 remain in the routing table.

upvoted 2 times

LanreDipeolu 3 months, 2 weeks ago

Selected Answer: C

C is the answer because Prefix-list goes with distribution-list not with access-list.

upvoted 1 times

jojoseb 5 months, 2 weeks ago

agree with B

upvoted 1 times

guy276465281819372 6 months, 3 weeks ago

Selected Answer: B

answer is B

upvoted 1 times

Malasxd 7 months, 2 weeks ago

Selected Answer: B

I'm sure it's B

upvoted 1 times

Dacusai 7 months, 3 weeks ago

I lab it and B is the correct one.

upvoted 1 times

anonymous1966 7 months, 4 weeks ago

Selected Answer: B

Confirmed now in PNET Lab.
Correct (B)

upvoted 1 times

davdtech 7 months, 4 weeks ago

We use a prefix list as it's name implies to match a list of subnets. In this case we only want to deny just one subnet. Now also in the question it does not specify if all other networks need to be denied. I go for B

upvoted 1 times

KingIT_ENG 8 months, 3 weeks ago

B is correct answer

upvoted 1 times

Wooker 9 months, 1 week ago

Selected Answer: B

B is correct!

upvoted 1 times

ArlindoCCNP 9 months, 1 week ago

B is correct!

upvoted 1 times

cchcano 11 months, 1 week ago

the route-map has an implied deny so if option B is chosen, a permit 20 should be added to the route-map to let the other routes pass. the correct option is the "C" ip prefix-list deny 192.168.2.2/32 with this you deny that route and allow the others, when it is evaluated in the route-map only the other routes arrive and as it starts with a permit all other routes pass with the exception of 192.168.2.2

upvoted 2 times

nicoaburto 11 months, 2 weeks ago

Selected Answer: C

the prefix-list is more relevant for filter with route-map

upvoted 2 times

Hurk2 11 months, 2 weeks ago

B is correct

upvoted 1 times

Question #9 Topic 1

What is a prerequisite for configuring BFD?

A. Jumbo frame support must be configured on the router that is using BFD.

B. All routers in the path between two BFD endpoints must have BFD enabled.

C. Cisco Express Forwarding must be enabled on all participating BFD endpoints.

D. To use BFD with BGP, the timers 3 9 command must first be configured in the BGP routing process.

Correct Answer: *C*

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fs_bfd.html#wp1043332

Community vote distribution

C (100%)

Ll123123 2 months ago

Selected Answer: C

choose C

upvoted 2 times

Alexloh 1 year, 6 months ago

Selected Answer: C

Agreed for C

upvoted 1 times

xziomal9 1 year, 8 months ago

The correct answer is: C

upvoted 1 times

Hack4 1 year, 10 months ago

the given answer is correct

upvoted 1 times

Girmiti 1 year, 11 months ago

Prerequisites for Bidirectional Forwarding Detection
•Cisco Express Forwarding (CEF) and IP routing must be enabled on all participating routers.
https://www.cisco.com/en/US/docs/ios/12_4t/ip_route/configuration/guide/t_bfd.html#wp1043332

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Sajj_gabi 3 years ago

Prerequisites for Bidirectional Forwarding Detection
Cisco Express Forwarding and IP routing must be enabled on all participating routers

upvoted 1 times

Guitarman 3 years, 4 months ago

CCIEBYDEC is correct, it's the very first pre-requisite.....it's C

upvoted 1 times

CCIEBYDEC 3 years, 4 months ago

answer is C https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fs_bfd.html#wp1043332

upvoted 2 times

Isolated 3 years, 6 months ago

Prerequisites for Bidirectional Forwarding Detection
One of the IP routing protocols supported by BFD must be configured on the routers before BFD is deployed. ... The router must be running BFD Version 1. The BFD session type must be IPv4 single hop. BFD echo mode must be disabled for the session.

upvoted 1 times

Question #10 *Topic 1*

DRAG DROP -

Drag and drop the OSPF adjacency states from the left onto the correct descriptions on the right.

Select and Place:

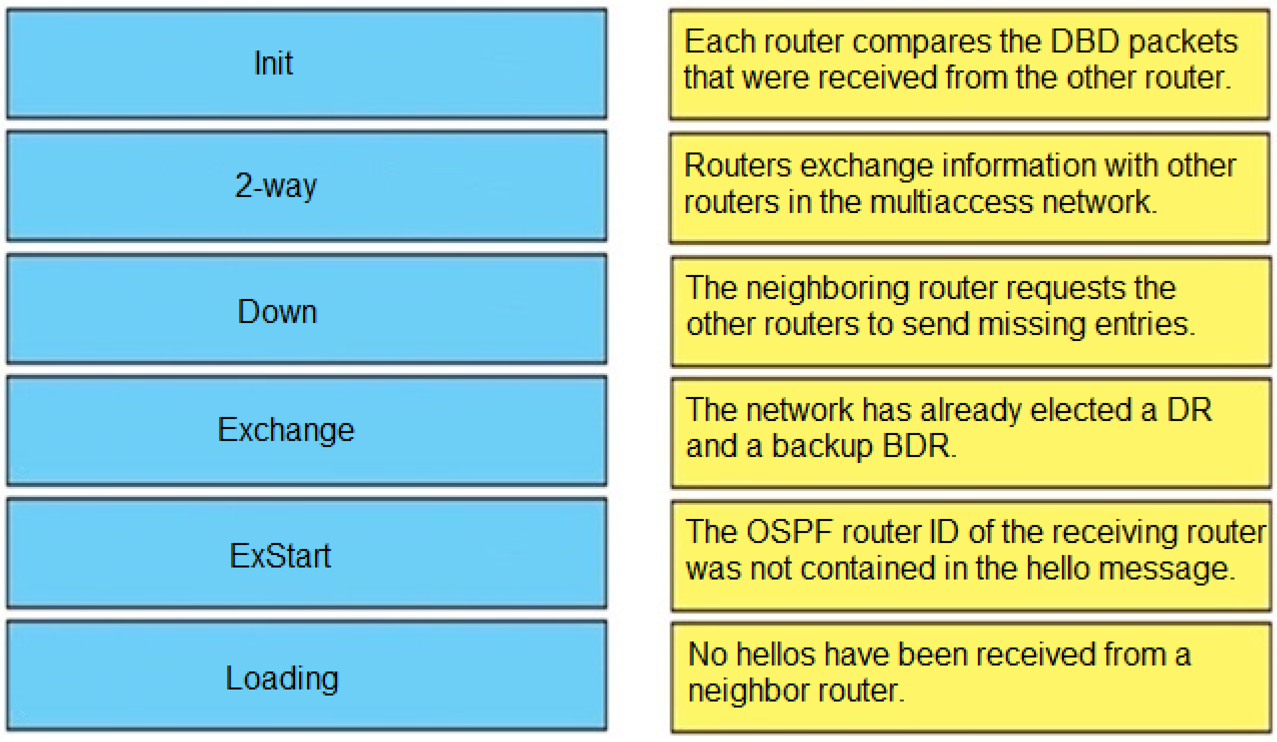

Init	Each router compares the DBD packets that were received from the other router.
2-way	Routers exchange information with other routers in the multiaccess network.
Down	The neighboring router requests the other routers to send missing entries.
Exchange	The network has already elected a DR and a backup BDR.
ExStart	The OSPF router ID of the receiving router was not contained in the hello message.
Loading	No hellos have been received from a neighbor router.

Correct Answer:

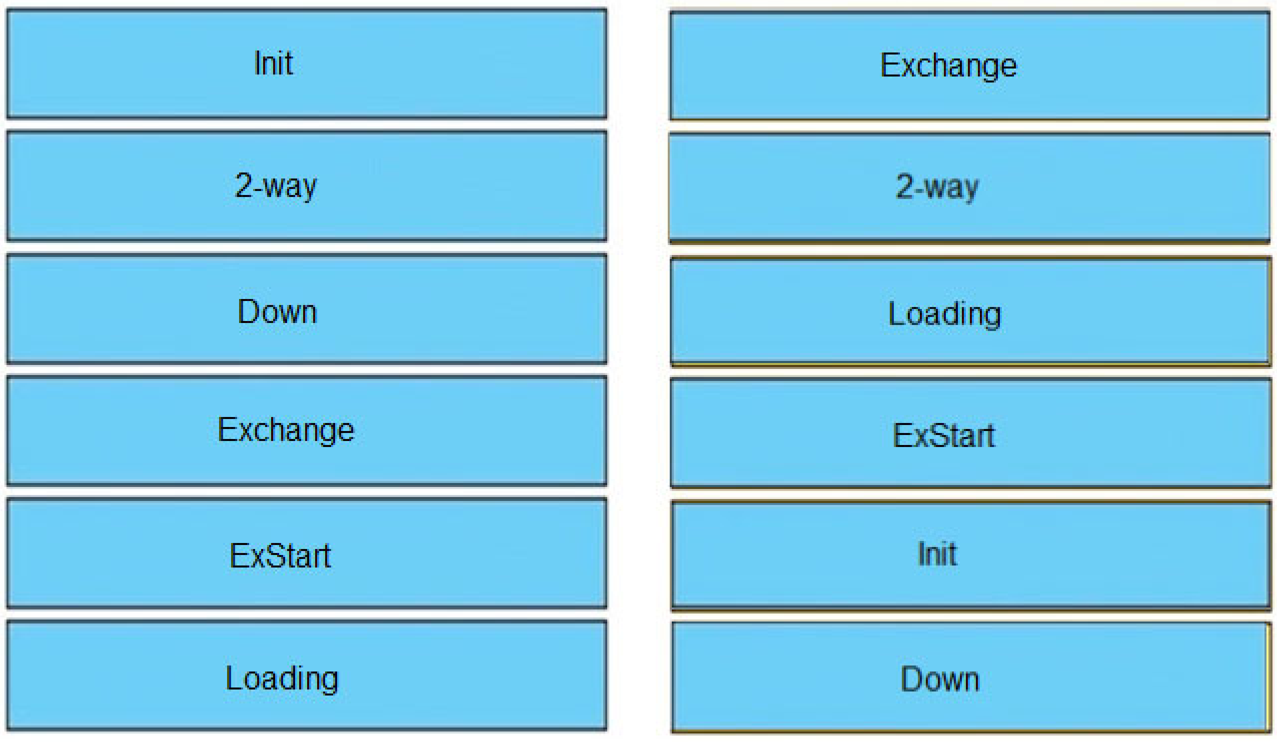

Init	Exchange
2-way	2-way
Down	Loading
Exchange	ExStart
ExStart	Init
Loading	Down

Reference:

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13685-13.html

Guitarman Highly Voted 3 years, 4 months ago

Are you guys sure about that? If you look, what you say shoulr be 2 way says that the DR and BDR have already been elected. The article referenced for ExStart says "Once the DR and BDR are elected, the actual process of exchanging link state information can start between the routers." That to me suggests that the DR and BDR have already been elected.

upvoted 7 times

bjromero28 Highly Voted 2 years, 2 months ago

1) Exchange - Routers exchange database descriptor (DBD) packets. Contents of the DBD received are compared to the information contained in the routers link-state database.

2) 2-Way - Each router has seen the other's hello packet. At the end of this stage, the DR and BDR for broadcast and non-broadcast multiacess networks are elected.

3) Loading - The actual exchange of link state information occurs. If a router receives an outdated or missing LSA, it requests that LSA by sending a link-state request packet.

4) Exstart - The routers and their DR and BDR establish a master-slave relationship.

5) Init - Specifies that the router has received a hello packet from neighbor, but receiving router's ID was not included in the hello packet.

6) Down - No Hellos have been received

Given Answer is correct.

Link: https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13685-13.html#intro

upvoted 6 times

HungarianDish 7 months ago

I agree on this solution.

upvoted 4 times

Remsync Most Recent 1 year, 2 months ago

The description for 2-way and Exstart is horrible.

upvoted 1 times

timtgh 1 year, 7 months ago

If this order is correct, their description of the 2-way state is awful.

upvoted 2 times

xziomal9 1 year, 8 months ago

Given answer is correct.

upvoted 1 times

studybuddy10 2 years, 1 month ago

Given answer is correct

upvoted 1 times

beatido 2 years, 3 months ago

Down and Loading need to be swapped around

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

frzzt123 2 years, 7 months ago

Unless answer was corrected in the meantime, at this very moment it is correct the way it is.
Nothing should be swapped IMO

upvoted 1 times

RHK0783 2 years, 8 months ago

The given answer is correct:
http://www.firewall.cx/networking-topics/routing/ospf-routing-protocol/1142-ospf-adjacency-neighbor-states-forming-process.html

upvoted 1 times

CraigB83 3 years, 2 months ago

It sounds right to me, Exchange is where DBD are compared

upvoted 2 times

james4231 3 years, 2 months ago

exchange and 2 way should be swapped. Exchange information should be mentioning the exchange stage, which "choose the initial sequence number for adjacency formation"

upvoted 1 times

CCIEBYDEC 3 years, 4 months ago

true, 2 way and Exstart need to be swapped.

upvoted 1 times

anonymous1966 3 years, 5 months ago

2-Way and ExStart positions should be inverted.

2-Way: "At the end of this stage, the DR and BDR for broadcast and non-broadcast multiacess networks are elected"
ExStart: "Once the DR and BDR are elected, the actual process of exchanging link state information can start between the routers and their DR and BDR"
Link: https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13685-13.html

upvoted 2 times

thissiteisgreat 2 years, 11 months ago

no, it is the 2-way state where DR / BDR is selected, once they are elected, it beings the next state, which is Exstart. 2-way state is the final state for DROTHERS where bidirectional communication has been established.

You misunderstood the reference in your link

upvoted 2 times

Refer to the exhibit. R2 is a route reflector, and R1 and R3 are route reflector clients. The route reflector learns the route to 172.16.25.0/24 from R1, but it does not advertise to R3.

What is the reason the route is not advertised?

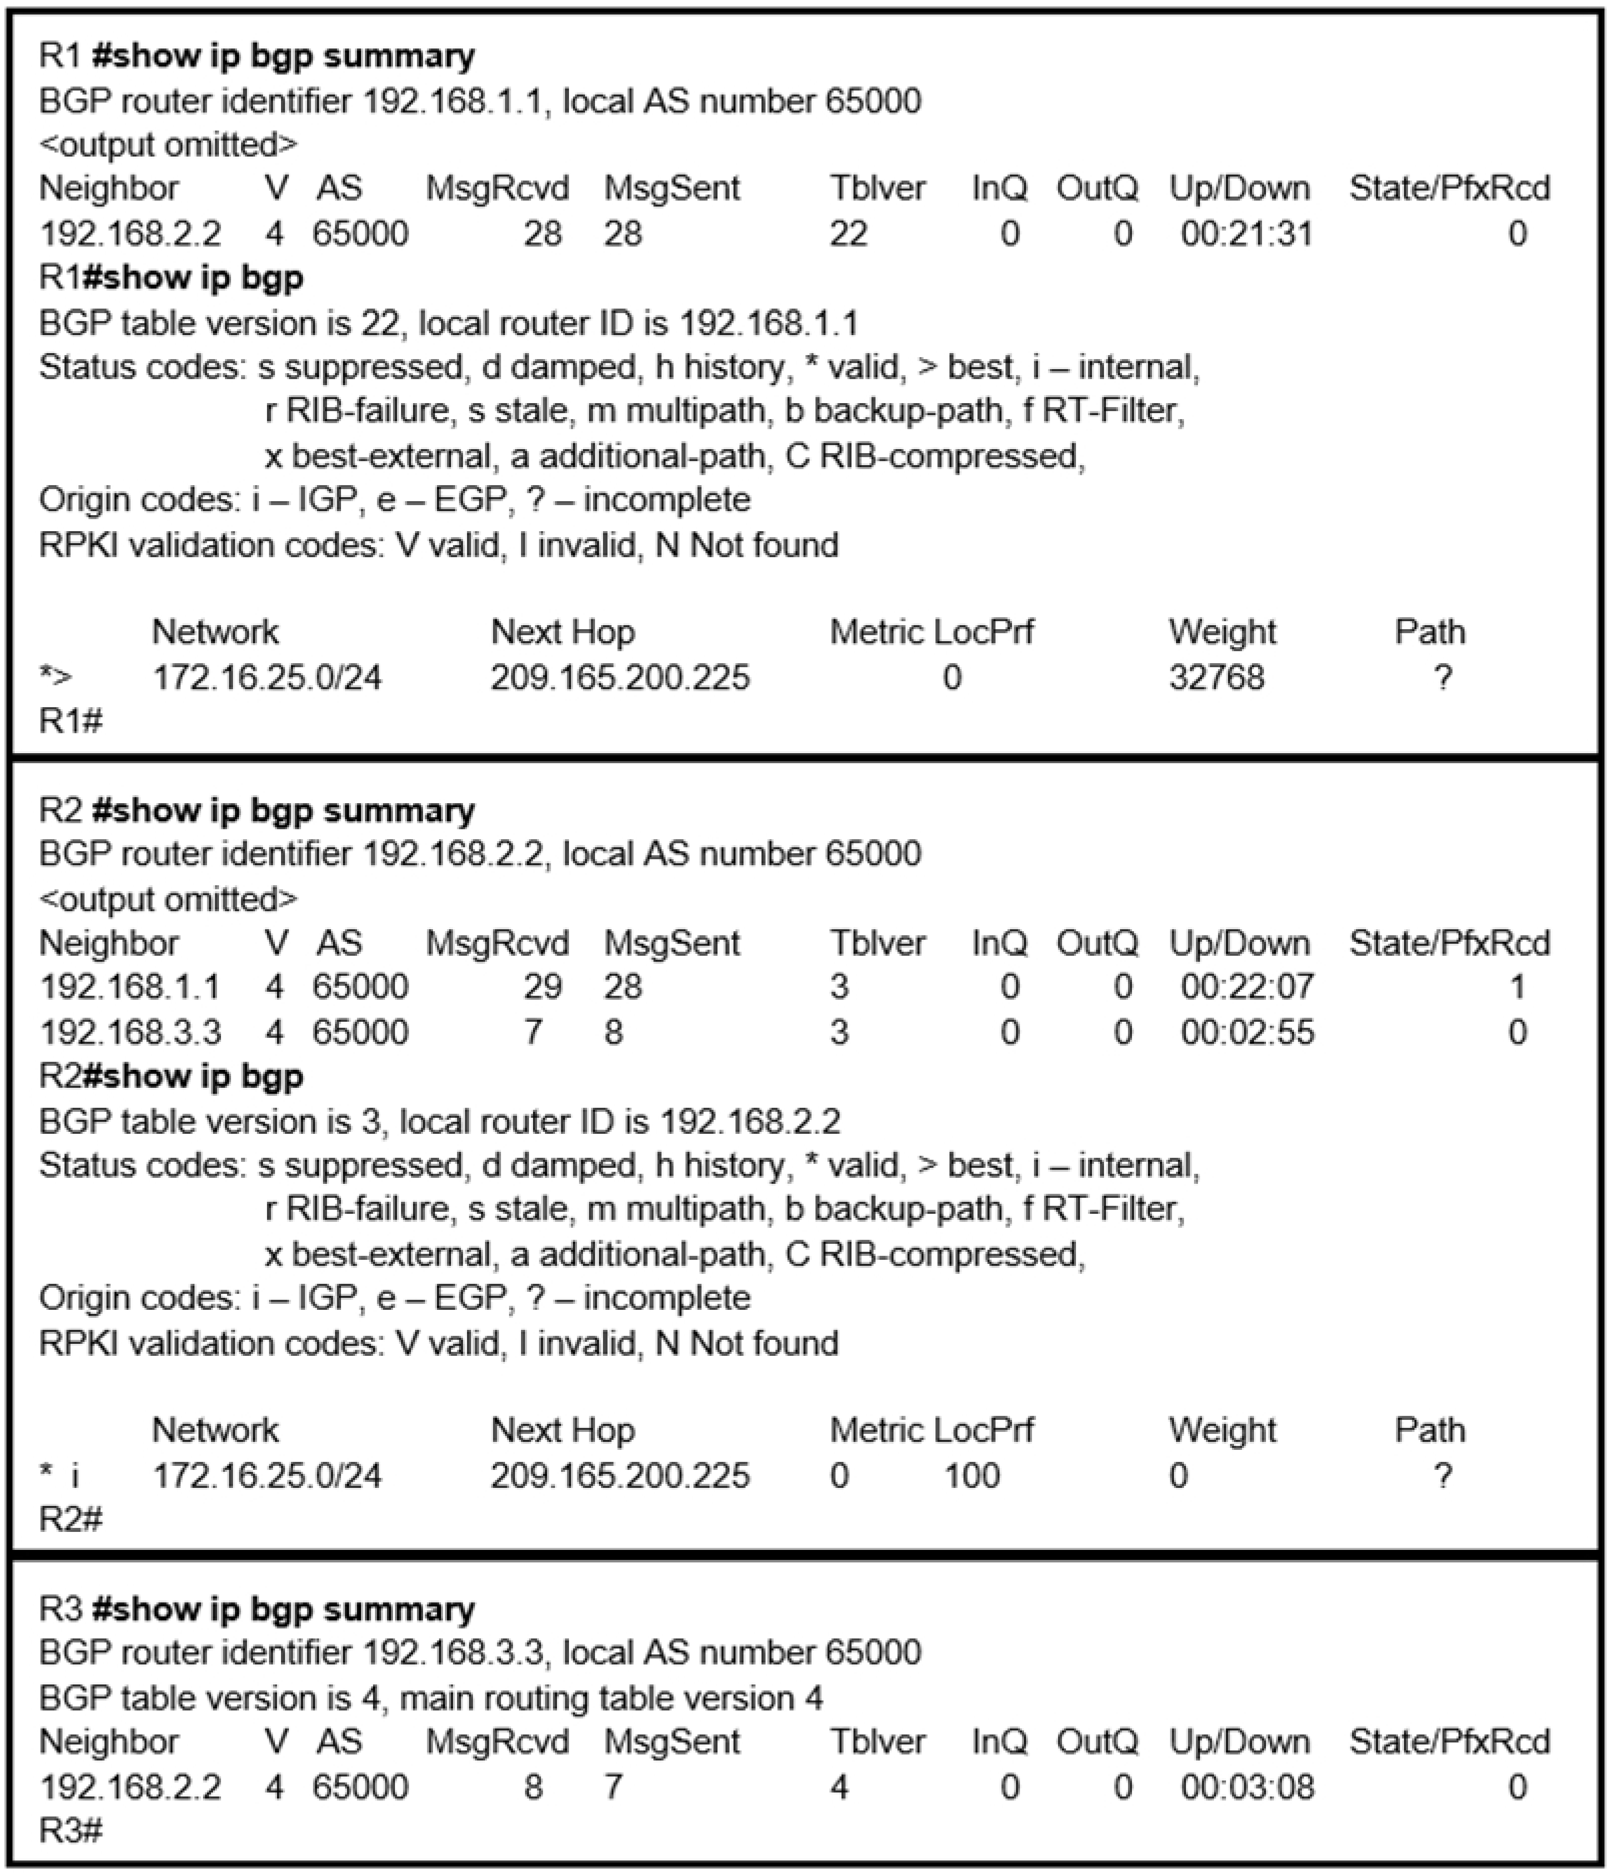

```
R1 #show ip bgp summary
BGP router identifier 192.168.1.1, local AS number 65000
<output omitted>
Neighbor      V  AS      MsgRcvd  MsgSent    Tblver   InQ  OutQ  Up/Down   State/PfxRcd
192.168.2.2   4  65000       28   28         22         0     0   00:21:31             0
R1#show ip bgp
BGP table version is 22, local router ID is 192.168.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i – internal,
              r RIB-failure, s stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, C RIB-compressed,
Origin codes: i – IGP, e – EGP, ? – incomplete
RPKI validation codes: V valid, I invalid, N Not found

        Network           Next Hop          Metric LocPrf     Weight     Path
*>      172.16.25.0/24    209.165.200.225         0           32768        ?
R1#
```

```
R2 #show ip bgp summary
BGP router identifier 192.168.2.2, local AS number 65000
<output omitted>
Neighbor      V  AS      MsgRcvd  MsgSent    Tblver   InQ  OutQ  Up/Down   State/PfxRcd
192.168.1.1   4  65000       29   28         3          0     0   00:22:07             1
192.168.3.3   4  65000        7   8          3          0     0   00:02:55             0
R2#show ip bgp
BGP table version is 3, local router ID is 192.168.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i – internal,
              r RIB-failure, s stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, C RIB-compressed,
Origin codes: i – IGP, e – EGP, ? – incomplete
RPKI validation codes: V valid, I invalid, N Not found

        Network           Next Hop          Metric LocPrf     Weight     Path
* i     172.16.25.0/24    209.165.200.225   0      100        0            ?
R2#
```

```
R3 #show ip bgp summary
BGP router identifier 192.168.3.3, local AS number 65000
BGP table version is 4, main routing table version 4
Neighbor      V  AS      MsgRcvd  MsgSent    Tblver   InQ  OutQ  Up/Down   State/PfxRcd
192.168.2.2   4  65000        8   7          4          0     0   00:03:08             0
R3#
```

A. R2 does not have a route to the next hop, so R2 does not advertise the prefix to other clients.

B. Route reflector setup requires full IBGP mesh between the routers.

C. In route reflector setup, only classful prefixes are advertised to other clients.

D. In route reflector setups, prefixes are not advertised from one client to another.

Correct Answer: *A*

Community vote distribution

A (100%)

⊟ **vdsdrs** Highly Voted 2 years, 4 months ago

Answer A is correct. You can see that on R2 route is missing '>' what means it's only in BGP table and not in RIB --> it will not be advertised

upvoted 8 times

⊟ **wts** 1 year, 3 months ago

Option A has a different reason.
upvoted 1 times

Malasxd Most Recent 7 months, 2 weeks ago

Selected Answer: A

The correct answer is: A
upvoted 1 times

Nhan 1 year, 3 months ago

The next hop on R1 and R2 is the same, R2 doesn't have the next hope to R3?
upvoted 1 times

xziomal9 1 year, 8 months ago

Selected Answer: A

The correct answer is: A
upvoted 1 times

Hack4 1 year, 10 months ago

A is correct
upvoted 1 times

Hack4 1 year, 10 months ago

A is correct
upvoted 1 times

Networkingguy 1 year, 10 months ago

Selected Answer: A

A looks to be correct here
upvoted 2 times

[Removed] 1 year, 11 months ago

Doesnt the * mean the route is valid indicating there is a next hop address??
upvoted 2 times

Jenia1 2 years ago

Will go for D, this is the closest answer, route reflector (R2) is receiving the route from R1, BGP between R2-R3 are established. I believe the command "neighbor 192.168.3.3 route-reflector-client " is missing on the R2 (not shown on the output), so R3 is not a client, only a BGP peer, so according to iBGP rules the R2 will not advertise the route that is received via IBG to non reflector clients
upvoted 1 times

Jenia1 2 years ago

Disregard my previous comment. Answer A is correct.
upvoted 3 times

wts 1 year, 3 months ago

Why A?
upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct
upvoted 3 times

AliMo123 2 years, 4 months ago

None of them is true
the topology is missing IGP routing protocol that's why R2 does not know how to reach the next hop, the closest answer is D
upvoted 1 times

[Removed] 1 year, 11 months ago

D makes absolutely no sense considering R3 is a client and R2 is a reflector...
upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct
upvoted 2 times

Chris_Li 2 years, 5 months ago

i think none of these 4 options is right...who knows which one is correct
upvoted 1 times

Benzzyy 2 years, 10 months ago

A is correct
upvoted 2 times

tomasz 2 years, 10 months ago

A is correct

upvoted 3 times

geek1992 2 years, 12 months ago

I'm going with A

upvoted 4 times

Sajj_gabi 3 years ago

your correct in what you say however- that's why RR is used to reflect the route to the other peer, IBGP peering can be configured so that it reflects routes to another IBGP peer.

upvoted 3 times

Question #12 *Topic 1*

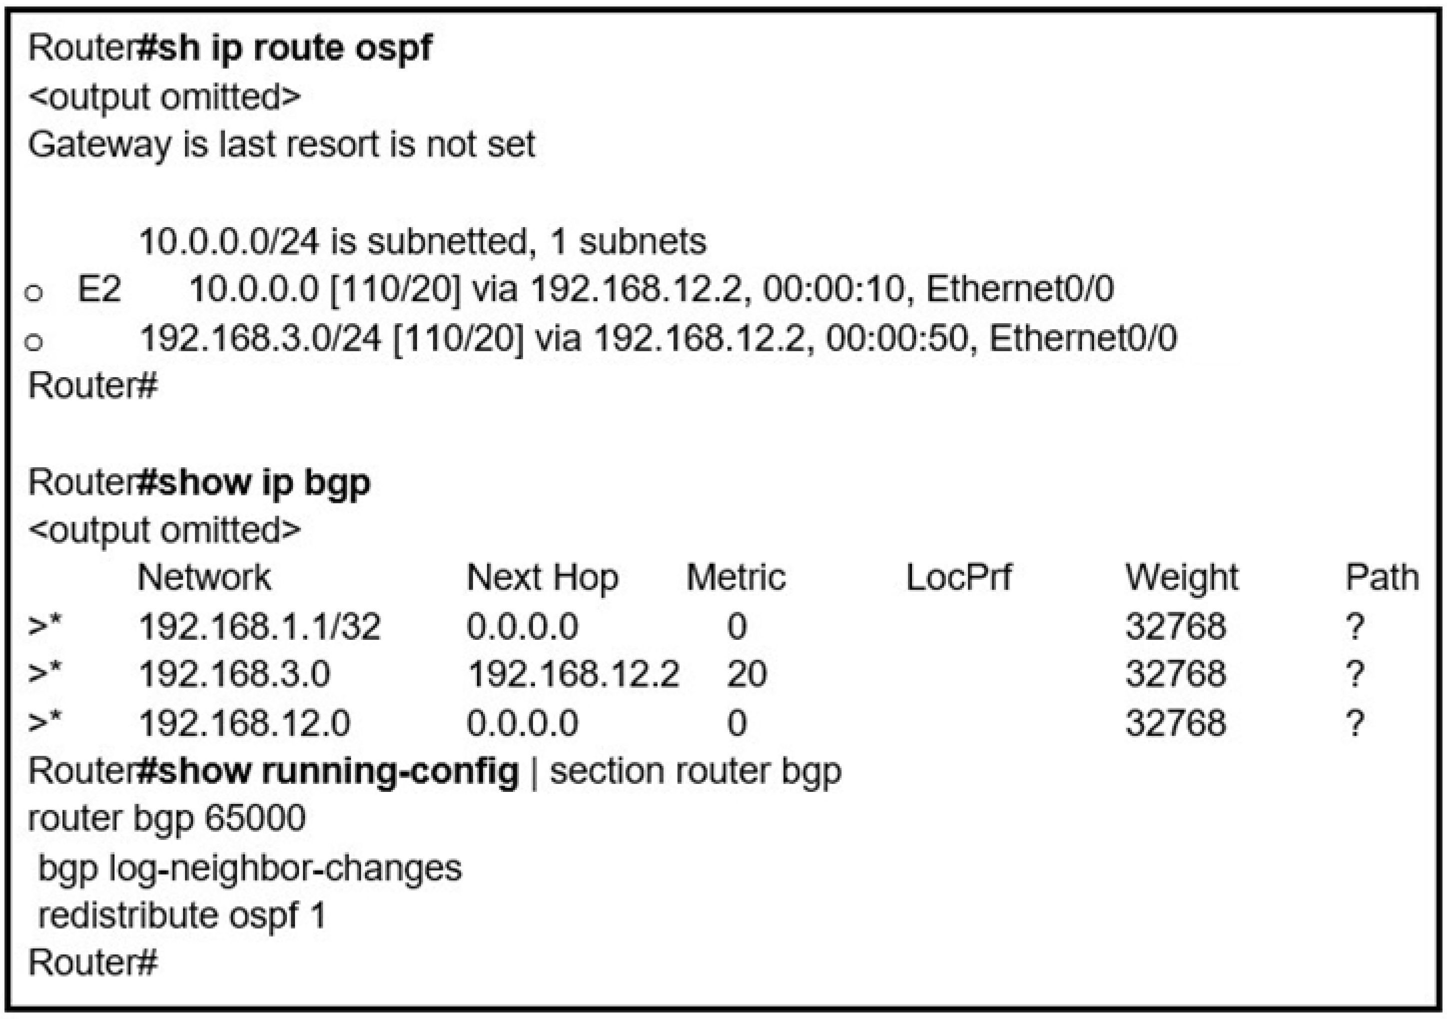

```
Router#sh ip route ospf
<output omitted>
Gateway is last resort is not set

          10.0.0.0/24 is subnetted, 1 subnets
o   E2      10.0.0.0 [110/20] via 192.168.12.2, 00:00:10, Ethernet0/0
o         192.168.3.0/24 [110/20] via 192.168.12.2, 00:00:50, Ethernet0/0
Router#

Router#show ip bgp
<output omitted>
          Network          Next Hop      Metric      LocPrf      Weight      Path
>*        192.168.1.1/32   0.0.0.0         0                     32768        ?
>*        192.168.3.0      192.168.12.2    20                    32768        ?
>*        192.168.12.0     0.0.0.0         0                     32768        ?
Router#show running-config | section router bgp
router bgp 65000
 bgp log-neighbor-changes
 redistribute ospf 1
Router#
```

Refer to the exhibit. An engineer is trying to redistribute OSPF to BGP, but not all of the routes are redistributed.

What is the reason for this issue?

A. By default, only internal routes and external type 1 routes are redistributed into BGP

B. Only classful networks are redistributed from OSPF to BGP

C. BGP convergence is slow, so the route will eventually be present in the BGP table

D. By default, only internal OSPF routes are redistributed into BGP

Correct Answer: *D*

Community vote distribution

D (100%)

heamgu Highly Voted 3 years, 6 months ago

The answer is correct is D.
Reference: https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redis.html?dtid=osscdc000283#redistributionofonlyospfinternalroutesintobgp

upvoted 8 times

Malasxd Most Recent 7 months, 2 weeks ago

Selected Answer: D

D is correct. If you configure the redistribution of OSPF into BGP without keywords, only OSPF intra-area and inter-area routes are redistributed into BGP, by default

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redis.html?dtid=osscdc000283#redistributionofonlyospfinternalroutesintobgp

upvoted 1 times

Remsync 1 year, 2 months ago

Selected Answer: D

D is correct.

"If you configure the redistribution of OSPF into BGP without keywords, only OSPF intra-area and inter-area routes are redistributed into BGP, by default. "

upvoted 3 times

tipama7298 1 year, 2 months ago

If you configure the redistribution of OSPF into BGP without keywords, only OSPF intra-area and inter-area routes are redistributed into BGP, by default. You can use the internal keyword along with the redistribute command under router bgp to redistribute OSPF intra- and inter-area routes.

upvoted 1 times

Router 1 year, 3 months ago

b is the correct ans, by default only classful network will be redistributed from ospf to other routing protocol unless you added subnet command at the end

upvoted 1 times

Remsync 1 year, 2 months ago

The "subnet" keyword is only used to redistribute INTO OSPF, not from.

https://learningnetwork.cisco.com/s/question/0D53i00000Kt6nCCAR/redistribute-subnet-keyword

upvoted 3 times

jarz 1 year, 4 months ago

After reading from the links provided to Cisco regarding redistributing OSPF into BGP, quoting directly from Cisco
Note: The configuration shows match external 1 external 2 and the command entered was redistribute ospf 1 match external. This is normal because OSPF automatically appends "external 1 external 2" in the configuration. It matches both OSPF external 1 and external 2 routes and it redistributes both routes into BGP.

So D is incorrect as well.

upvoted 2 times

Remsync 1 year, 2 months ago

But what you're quoting is on the section to, explicitly, redistribute Only OSPF External (type 1 and 2) into BGP.

On the section above, it talks about the redistribution of OSPF internal routes into BGP and it says that that is the default redistribution (with no keywords):

"If you configure the redistribution of OSPF into BGP without keywords, only OSPF intra-area and inter-area routes are redistributed into BGP, by default."

upvoted 1 times

Alexloh 1 year, 5 months ago

Selected Answer: D

The correct answer is D

upvoted 1 times

Nhan 1 year, 6 months ago

In this case the route was marked with E2 is the OSPF external router from another Area won't be redistributed

upvoted 1 times

xziomal9 1 year, 8 months ago

Selected Answer: D

The correct answer is: D

upvoted 1 times

Hack4 1 year, 10 months ago

The given answer is correct

upvoted 1 times

Networkingguy 1 year, 10 months ago

Selected Answer: D

D looks to be correct here

upvoted 1 times

gndrx78 2 years ago

D
Probably to avoid risking a loop advertising external routes outside OSPF domain that can cause a loop not detected by BGP due to lack of ASN in OSPF info during redistribution?

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

akbntc 3 years ago

D is correct.

upvoted 2 times

CCIEBYDEC 3 years, 4 months ago

Answer is D. https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospfredis.html

upvoted 4 times

Question #13 Topic 1

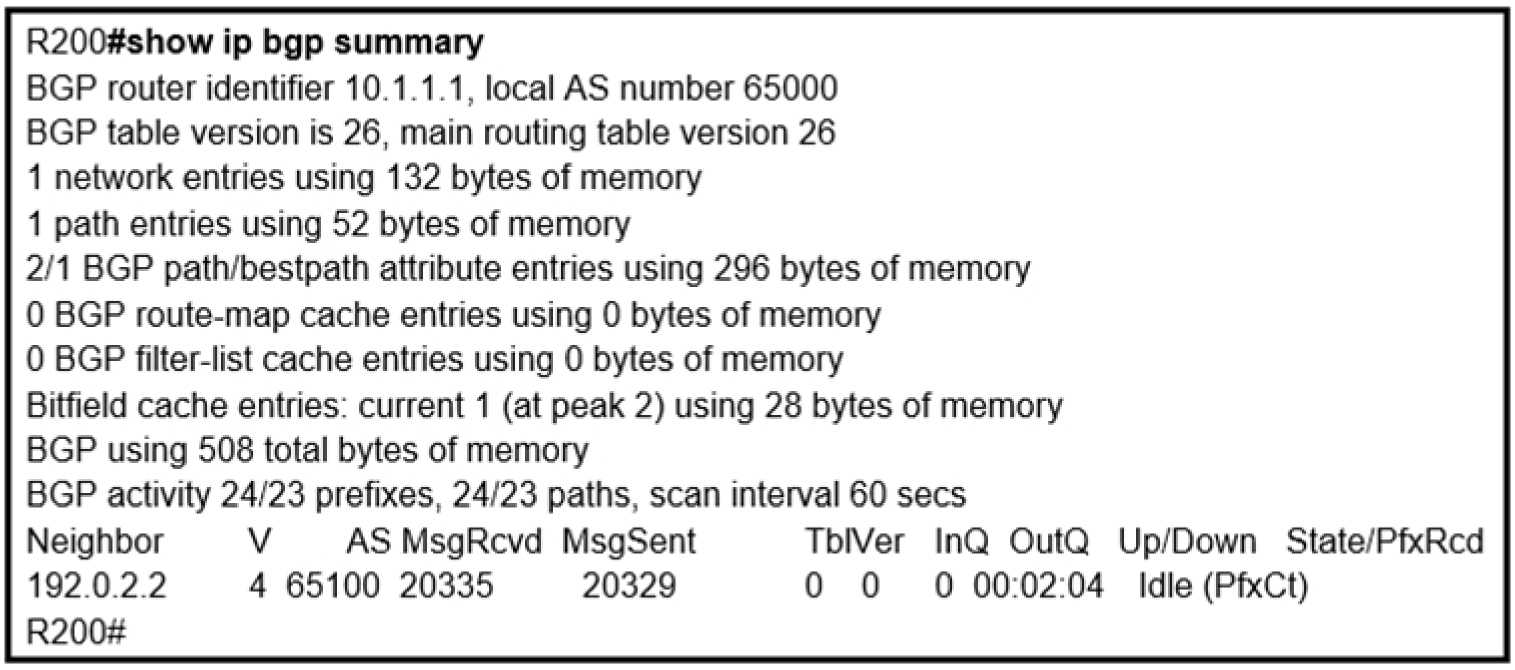

```
R200#show ip bgp summary
BGP router identifier 10.1.1.1, local AS number 65000
BGP table version is 26, main routing table version 26
1 network entries using 132 bytes of memory
1 path entries using 52 bytes of memory
2/1 BGP path/bestpath attribute entries using 296 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Bitfield cache entries: current 1 (at peak 2) using 28 bytes of memory
BGP using 508 total bytes of memory
BGP activity 24/23 prefixes, 24/23 paths, scan interval 60 secs
Neighbor      V    AS MsgRcvd MsgSent    TblVer  InQ OutQ Up/Down  State/PfxRcd
192.0.2.2     4 65100  20335   20329      0    0    0 00:02:04  Idle (PfxCt)
R200#
```

Refer to the exhibit. In which circumstance does the BGP neighbor remain in the idle condition?

A. if prefixes are not received from the BGP peer

B. if prefixes reach the maximum limit

C. if a prefix list is applied on the inbound direction

D. if prefixes exceed the maximum limit

Correct Answer: *D*

Community vote distribution

D (100%)

CraigB83 Highly Voted 3 years, 2 months ago

D
"The BGP Maximum-Prefix feature allows you to control how many prefixes can be received from a neighbor. By default, this feature allows a router to bring down a peer when the number of received prefixes from that peer exceeds the configured Maximum-Prefix limit"

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/25160-bgp-maximum-prefix.html#:~:text=The%20BGP%20Maximum%2DPrefix%20feature%20allows%20you%20to%20control%20how,the%20configured%20Maximum%2DPrefix%20limit.

upvoted 7 times

Girmiti Highly Voted 1 year, 11 months ago

D is the Answer
Idle (PfxCt) means the session is in the Idle state because the neighbor has sent more prefixes than the configured maximum-prefixes limit.

upvoted 5 times

Alexloh Most Recent 1 year, 5 months ago

Selected Answer: D

The correct answer is D

upvoted 2 times

Reikidude00 1 year, 6 months ago

how we can understand that maximum-prefix is being configured based on this output?

upvoted 1 times

xziomal9 1 year, 8 months ago

Selected Answer: D

The correct answer is: D

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 3 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

⊟ **thissiteisgreat** 2 years, 11 months ago

D is correct because there is the "PfxRcd" string under the State/PfxRcd field.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-s1.html#wp1583714062

upvoted 3 times

⊟ **Jack1188** 3 years, 4 months ago

D is the correct once.

upvoted 2 times

Question #14 Topic 1

Which attribute eliminates LFAs that belong to protected paths in situations where links in a network are connected through a common fiber?

A. shared risk link group-disjoint

B. linecard-disjoint

C. lowest-repair-path-metric

D. interface-disjoint

Correct Answer: *B*

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/xe-3s/asr1000/ire-xe-3s-asr1000/ire-ipfrr.html

Community vote distribution

A (100%)

Stupid Highly Voted 1 year, 11 months ago

Selected Answer: A

A seems to be right, https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/xe-3s/asr1000/ire-xe-3s-asr1000/ire-ipfrr.html#:~:text=Shared%20Risk%20Link,group%20share%20risks.

upvoted 6 times

Cisco_TechniciaN Most Recent 3 months, 3 weeks ago

Selected Answer: A

SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may also fail. Therefore, links in a group share risks.

upvoted 1 times

goomisch 8 months, 1 week ago

A is correct - Shared Risk Link Group (SRLG)-disjoint—Eliminates LFAs that belong to any of the protected path SRLGs. SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may also fail. Therefore, links in a group share risks.

upvoted 2 times

SDWAN 1 year, 2 months ago

appeared in my exam, along with several DNA questions that really shouldn't be here!

upvoted 2 times

jarz 1 year, 2 months ago

Has this question appeared in anyone's exam?

upvoted 2 times

networkWiz 1 year, 4 months ago

Selected Answer: A

LFA Tie-Breaking Rules

• Shared Risk Link Group (SRLG)-disjoint—Eliminates LFAs that belong to any of the protected path SRLGs.SRLGsrefer to situations where linksin a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may also fail. Therefore, links in a group share risks.

ref: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/xe-3s/asr1000/ire-xe-3s-asr1000.pdf

upvoted 2 times

Alexloh 1 year, 5 months ago

Selected Answer: A

Eliminates LFAs that belong to any of the protected path SRLGs. SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may also fail. Therefore, links in a group share risks.

The answer is A.

upvoted 1 times

Darcy42 1 year, 6 months ago

A is correct

upvoted 1 times

xziomal9 1 year, 8 months ago

Selected Answer: A

The correct answer is: A

upvoted 1 times

YaPet 1 year, 10 months ago

Selected Answer: A

I agree that A is true

upvoted 2 times

Networkingguy 1 year, 10 months ago

Selected Answer: A

A looks to be correct here

upvoted 1 times

yoyo_simon 2 years, 3 months ago

should be A correct

upvoted 2 times

examShark 2 years, 4 months ago

A is the correct answer

upvoted 2 times

tcze 2 years, 4 months ago

Correct Answer is A : Shared Risk Link Group (SRLG)-disjoint

Source : https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/xe-3s/asr1000/ire-xe-3s-asr1000/ire-ipfrr.html

upvoted 2 times

mynamelukecisco 2 years, 6 months ago

Shared Risk Link Group (SRLG)-disjoint—Eliminates LFAs that belong to any of the protected path SRLGs. SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may also fail. Therefore, links in a group share risks.

upvoted 2 times

ZachTL11 2 years, 8 months ago

A - Shared Risk Link Group (SRLG)-disjoint—Eliminates LFAs that belong to any of the protected path SRLGs. SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may also fail. Therefore, links in a group share risks.

upvoted 2 times

RHK0783 2 years, 8 months ago

Interface-disjoint—Eliminates LFAs that share the outgoing interface with the protected path.

Linecard-disjoint—Eliminates LFAs that share the line card with the protected path.

Lowest-repair-path-metric—Eliminates LFAs whose metric to the protected prefix is high. Multiple LFAs with the same lowest path metric may remain in the routing table after this tie-breaker is applied.

Shared Risk Link Group (SRLG)-disjoint—Eliminates LFAs that belong to any of the protected path SRLGs. SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may also fail. Therefore, links in a group share risks.

upvoted 3 times

Question #15 *Topic 1*

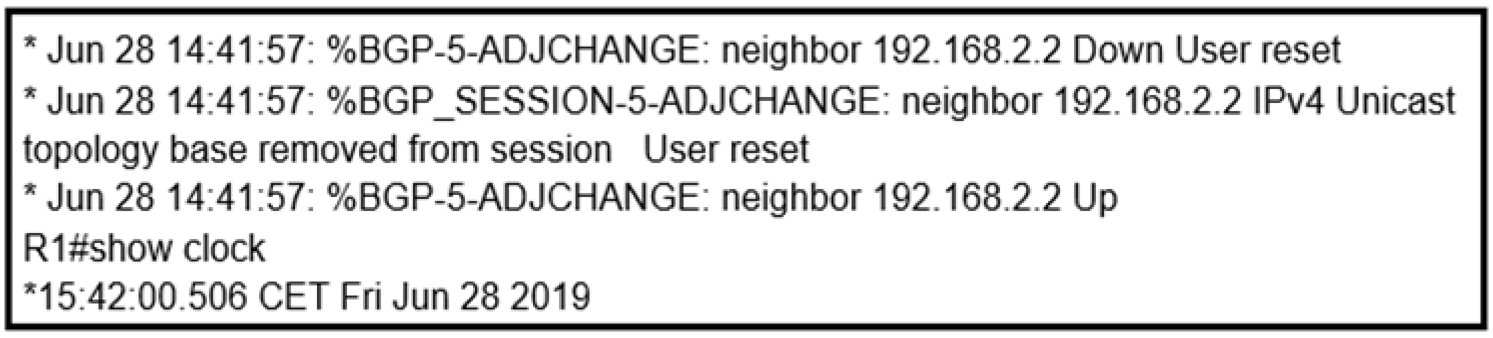

```
* Jun 28 14:41:57: %BGP-5-ADJCHANGE: neighbor 192.168.2.2 Down User reset
* Jun 28 14:41:57: %BGP_SESSION-5-ADJCHANGE: neighbor 192.168.2.2 IPv4 Unicast
topology base removed from session   User reset
* Jun 28 14:41:57: %BGP-5-ADJCHANGE: neighbor 192.168.2.2 Up
R1#show clock
*15:42:00.506 CET Fri Jun 28 2019
```

Refer to the exhibit. An engineer is troubleshooting BGP on a device but discovers that the clock on the device does not correspond to the time stamp of the log entries.
Which action ensures consistency between the two times?

A. Configure the service timestamps log uptime command in global configuration mode.

B. Configure the logging clock synchronize command in global configuration mode.

C. Configure the service timestamps log datetime localtime command in global configuration mode.

D. Make sure that the clock on the device is synchronized with an NTP server.

Correct Answer: *D*

Community vote distribution

C (80%) D (20%)

S_E_T Highly Voted 👍 3 years, 6 months ago

C is correct
https://community.cisco.com/t5/networking-documents/router-log-timestamp-entries-are-different-from-the-system-clock/ta-p/3132258

upvoted 10 times

Chiaretta Most Recent ⊙ 5 months, 3 weeks ago

Selected Answer: D

D is correct

upvoted 2 times

Malasxd 7 months, 2 weeks ago

It does not say the device timer is incorrect. It's says the device time and log time are different and you need to resolve it.

upvoted 1 times

Malasxd 7 months, 2 weeks ago

C is correct

upvoted 1 times

Koume 11 months, 1 week ago

Selected Answer: C

I vote C as the question is referring to the difference between log an the clock and this is fixed with service timestamp

upvoted 1 times

PrettyStupid 1 year ago

Selected Answer: C

I'm going with C
Reference: https://community.cisco.com/t5/networking-knowledge-base/router-log-timestamp-entries-are-different-from-the-system-clock/ta-p/3132258 and https://conetrix.com/blog/timestamps-on-logs-of-cisco-devices-do-not-match-actual-time-on-device

upvoted 1 times

Nhan 1 year, 3 months ago

Hey SET thank you for the link, you are the man, C is correct answer. Again thank you buddy

upvoted 1 times

networkWiz 1 year, 4 months ago

Selected Answer: C

C is the correct answer

upvoted 1 times

Pbshah 1 year, 5 months ago

Selected Answer: C

Even we synchronize the clock but it may show different timezone so we should set the “localtime” keyword (which uses local time zone for timestamps) so that the time of logging messages is matched with our clock.

upvoted 1 times

Alexloh 1 year, 5 months ago

Selected Answer: D

The answer is D

upvoted 1 times

Nhan 1 year, 6 months ago

C and D are both correct answer for this scenario, i would like to go with D. NTP server provide much more accurate clock setting than local device clock.

upvoted 3 times

xziomal9 1 year, 8 months ago

Selected Answer: C

The correct answer is: C

upvoted 1 times

Nhan 1 year, 8 months ago

D is the best answer, manually configure the clock is never can be as accurate as NTP server.

upvoted 1 times

davdtech 1 year, 10 months ago

I stick to D
There is an asterisks in front of the time meaning that the device is not in sync with an NTP

upvoted 3 times

Nonono 1 year, 10 months ago

C is correct

upvoted 1 times

Hack4 1 year, 10 months ago

C is the right answer

upvoted 2 times

Networkingguy 1 year, 10 months ago

Selected Answer: C

C looks to be correct here

upvoted 1 times

Jenia1 1 year, 11 months ago

Selected Answer: C

The correct answer is C

upvoted 2 times

Question #16 *Topic 1*

Refer to the exhibit. What is the result of applying this configuration?

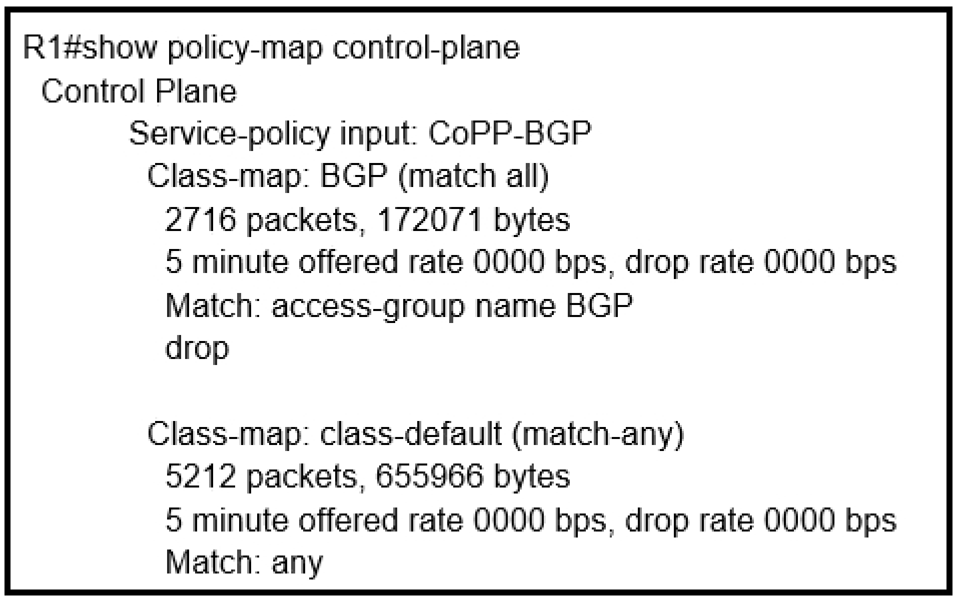

```
R1#show policy-map control-plane
 Control Plane
      Service-policy input: CoPP-BGP
       Class-map: BGP (match all)
        2716 packets, 172071 bytes
        5 minute offered rate 0000 bps, drop rate 0000 bps
        Match: access-group name BGP
        drop

       Class-map: class-default (match-any)
        5212 packets, 655966 bytes
        5 minute offered rate 0000 bps, drop rate 0000 bps
        Match: any
```

A. The router can form BGP neighborships with any other device.

B. The router cannot form BGP neighborships with any other device.

C. The router cannot form BGP neighborships with any device that is matched by the access list named ג€BGPג€.

D. The router can form BGP neighborships with any device that is matched by the access list named ג€BGPג€.

Correct Answer: *A*

Community vote distribution

C (100%)

Koume 11 months, 1 week ago

Selected Answer: C

Labbed with 3 routers
the peers i set with ACL in class map could no establish session so C is correct.

upvoted 3 times

Hurk2 11 months, 2 weeks ago

Selected Answer: C

C is correct

upvoted 3 times

Zizu007 11 months, 2 weeks ago

Selected Answer: C

with this (below) ACL both incoming (179) and outgoing (179) are blocked. BGP cannot be established.

```
R7#show ip access-lists
Extended IP access list ACL_BGP
10 permit tcp any any eq bgp (45 matches)
20 permit tcp any eq bgp any (3 matches)

R7#sh policy-map control-plane
Control Plane

Service-policy input: CoPP_IN

Class-map: CL_BGP (match-all)
76 packets, 4826 bytes
5 minute offered rate 0000 bps, drop rate 0000 bps
Match: access-group name ACL_BGP
drop

Class-map: class-default (match-any)
237 packets, 55172 bytes
5 minute offered rate 0000 bps, drop rate 0000 bps
Match: any
R7#
```

upvoted 1 times

kaisehhop 1 year, 1 month ago

Selected Answer: C

The correct answer is C

upvoted 2 times

Alexloh 1 year, 5 months ago

Selected Answer: C

The correct answer is C

upvoted 1 times

davdtech 1 year, 6 months ago

Ok so if the router can form BGP neighbourships with any other device, what are the marked packets 2716 ? These are dropped packets no ?

upvoted 1 times

zzmejce 1 year, 7 months ago

Selected Answer: C

The correct answer is: C

upvoted 1 times

xziomal9 1 year, 8 months ago

Selected Answer: C

The correct answer is: C

upvoted 1 times

Hack4 1 year, 10 months ago

The given answer is correct then A. The question refers about the control-plane protection mechanism.. The configuration shows that the router is still gonna etablish the BGP relationship to a given number of peers, but not all(because of policy assigned to that class-map based on rate-limit condition)

upvoted 3 times

Networkingguy 1 year, 10 months ago

Selected Answer: C

Its C, whoever admins this site is a nuffie.

upvoted 2 times

[Removed] 1 year, 11 months ago

I dont see how it can be any answer other than C. A tcp connection is required for BGP adjacencies to form. When the responding router matching the BGP acl sends its response packet its going to get dropped...

upvoted 3 times

gndrx78 2 years ago

Selected Answer: C

C seems the most logical considered some packets have matched and some other not

upvoted 4 times

studybuddy10 2 years, 1 month ago

C - labbed and existing neighbours that matched the ACL go down.

upvoted 2 times

Raider1 2 years, 2 months ago

Not sure if the answer is A or C. One class-map states drop any thing name BGP, and another Class map states allow any.

upvoted 2 times

error_909 2 years, 2 months ago

The correct answer is C

upvoted 1 times

yoyo_simon 2 years, 3 months ago

C should be correct

upvoted 1 times

examShark 2 years, 4 months ago

The correct answer is C

upvoted 3 times

Question #17

Topic 1

Which command displays the IP routing table information that is associated with VRF-Lite?

A. show ip vrf

B. show ip route vrf

C. show run vrf

D. show ip protocols vrf

Correct Answer: *B*

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2/50sg/configuration/guide/Wrapper-46SG/vrf.html#wp1045708

Community vote distribution

B (100%)

Alexloh 1 year, 5 months ago

Selected Answer: B

The answer is B

upvoted 3 times

xziomal9 1 year, 8 months ago

Selected Answer: B

The correct answer is: B

upvoted 2 times

Girmiti 1 year, 11 months ago

Selected Answer: B

show ip route vrf (vrf-name)

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

Question #18 Topic 1

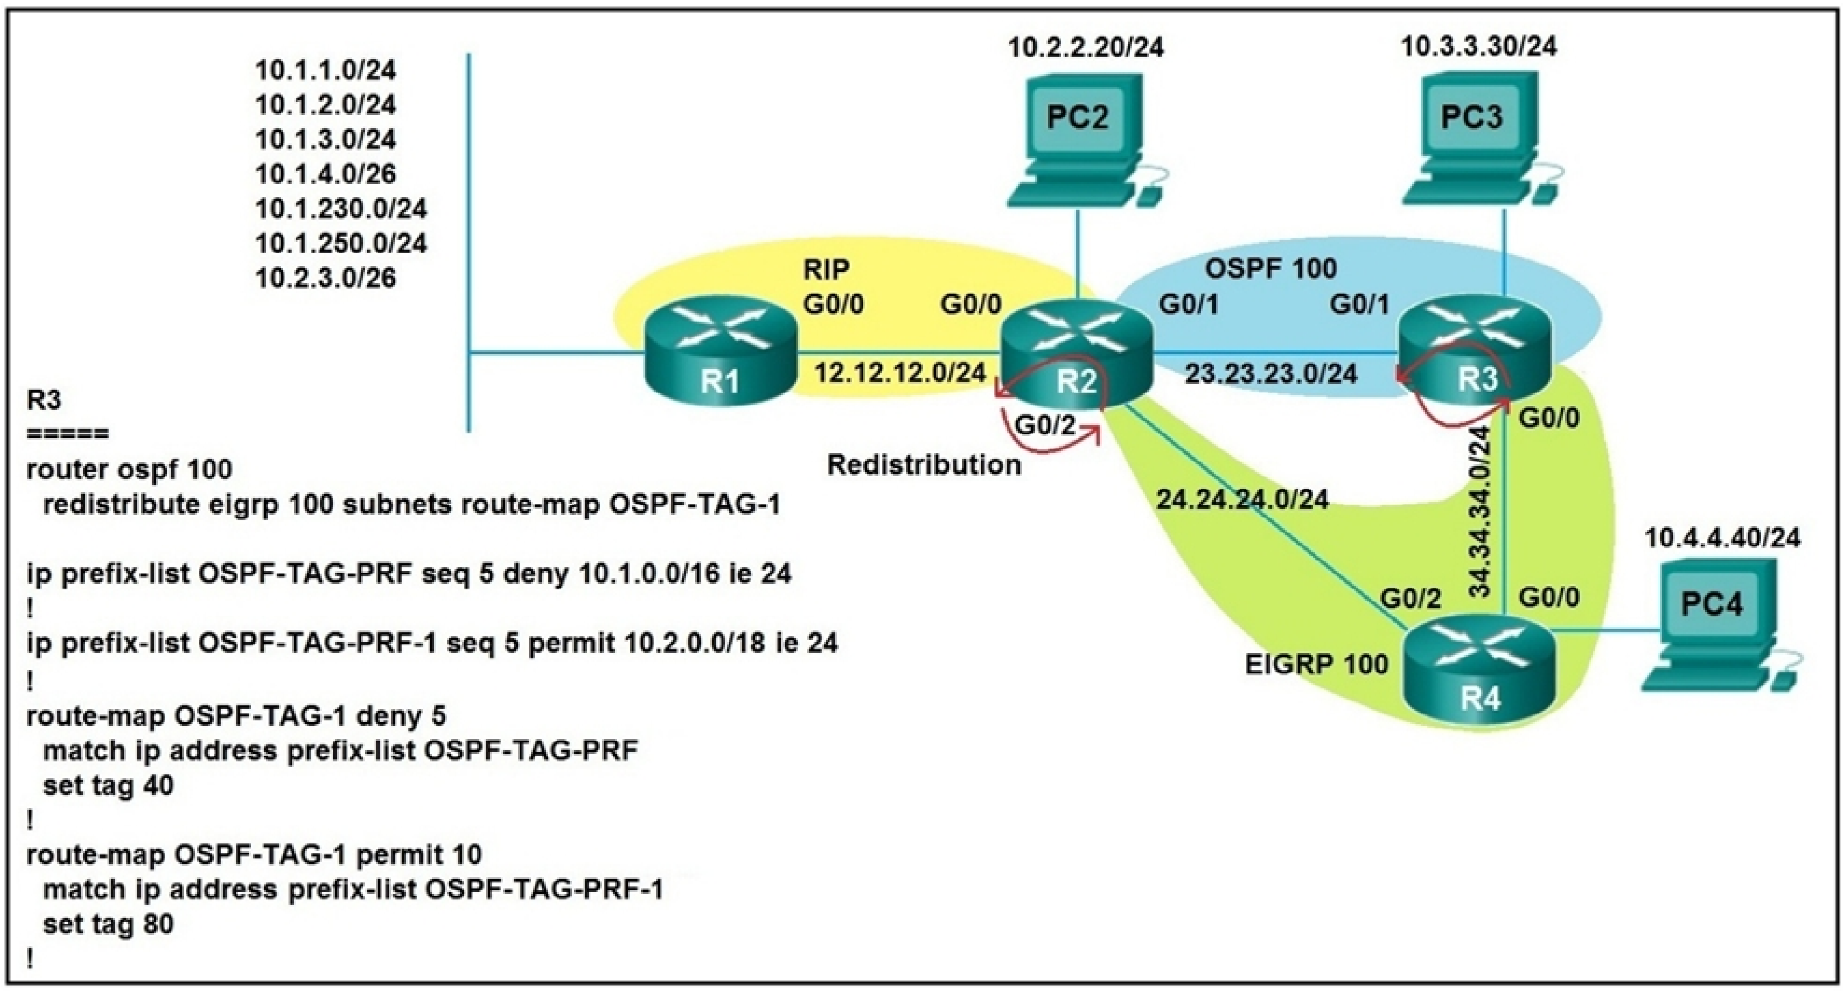

```
R3
=====
router ospf 100
 redistribute eigrp 100 subnets route-map OSPF-TAG-1

ip prefix-list OSPF-TAG-PRF seq 5 deny 10.1.0.0/16 ie 24
!
ip prefix-list OSPF-TAG-PRF-1 seq 5 permit 10.2.0.0/18 ie 24
!
route-map OSPF-TAG-1 deny 5
 match ip address prefix-list OSPF-TAG-PRF
 set tag 40
!
route-map OSPF-TAG-1 permit 10
 match ip address prefix-list OSPF-TAG-PRF-1
 set tag 80
!
```

Refer to the exhibit. Which subnet is redistributed from EIGRP to OSPF routing protocols?

A. 10.2.2.0/24

B. 10.1.4.0/26

C. 10.1.2.0/24

D. 10.2.3.0/26

Correct Answer: *A*

Community vote distribution

A (100%)

David98898998 6 months, 4 weeks ago

Selected Answer: A

Tested on GNS3. It's stupid as hell and I feel dumber for doing it, but A is definitely the answer.

upvoted 4 times

David98898998 6 months, 3 weeks ago

Route-maps will not match on deny ACLs or deny statement prefix-lists. They will ignore them. Sequence 5 of the route map is entirely ignored.

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

I agree. Prefixes from network 10.1.0.0/16 with length /16-24 are not evaluated in seq 5, but are denied by implicit deny-all at the end of the route-map.

upvoted 1 times

Iarn 1 year, 7 months ago

Bit confused why every has A given the logic shown for matching answer A 10.2.3.0/26 would also match?!

upvoted 2 times

Iarn 1 year, 7 months ago

On second thought le 24 is /0-24 thus /26 is greater

upvoted 3 times

xziomal9 1 year, 8 months ago

Selected Answer: A

The correct answer is: A

upvoted 1 times

thanh123 1 year, 8 months ago

I'm with A, too

upvoted 1 times

Networkingguy 1 year, 10 months ago

Selected Answer: A

A is correct here

upvoted 1 times

ciscomicha 1 year, 11 months ago

Selected Answer: A

I'm with A. Given answer. It is the only route that match an route-map permit statement because it matches the second prefix-list

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

beatido 2 years, 3 months ago

Its clearly A

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

RTE 2 years, 5 months ago

A is right, permit statement in second route-map and permit int prefix-list with network length <=24, implicit deny at the end of r-map

upvoted 2 times

azharken 2 years, 7 months ago

wrong question
both prefix lists are permitting

upvoted 2 times

ichweissauchnicht 2 years ago

That's true (deny in first acl and deny in route-map => permit). This question is strange...

upvoted 1 times

JOKERR 2 years ago

No. Deny in the ACL or Prefix list mean that entry is not affected by the route map. Deny means let the route pass. Permit means route map is permitted to take action on that entry.

upvoted 3 times

[Removed] 1 year, 11 months ago

Even if both are permitting there's a catch all class map(implicit deny) at the end and it will match that and be denied.

upvoted 1 times

oasc 2 years, 8 months ago

C is the one correct

upvoted 3 times

Pb1805 2 years, 7 months ago

What about A?

upvoted 3 times

Question #19 *Topic 1*

Which configuration adds an IPv4 interface to an OSPFv3 process in OSPFv3 address family configuration?

A. router ospfv3 1 address-family ipv4

B. Router(config-router)#ospfv3 1 ipv4 area 0

C. Router(config-if)#ospfv3 1 ipv4 area 0

D. router ospfv3 1 address-family ipv4 unicast

Correct Answer: *D*

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/xe-3s/iro-xe-3s-book/ip6-route-ospfv3-add-fam-xe.html

Community vote distribution

C (100%)

samne168 Highly Voted 3 years, 6 months ago

The correct answer C:
Device(config-if)# ospfv3 1 area 1 ipv4
because the question is which command add ipv4 interface to OSPFv3

upvoted 18 times

[Removed] 2 years, 1 month ago

The 2nd half of the question asks for the config under address-family. Once you create the process you enter address-family config where you would then enable ipv4 address-family. Thats why the answer is D. If they would have asked on a specific int. it would then be C.

upvoted 2 times

Jenia1 1 year, 10 months ago

But you can't add an interface using the address-family command in OSPFv3, as there is no network statement, whatever you will configure under the address-family will not take any effect until you add the interface using: ospfv3 1 area 1 ipv4 in interface configuration mode.
According to the question, C should be correct

upvoted 2 times

Brand Most Recent 3 months, 4 weeks ago

Selected Answer: C

```
R1(config-if)#ipv6 enable
R1(config-if)#ospfv3 1 ipv4 area 0
R1(config-if)#do show run | sec ospf
ospfv3 1 ipv4 area 0
router ospfv3 1
!
address-family ipv4 unicast
exit-address-family
```

Lab it people... It's "C"

upvoted 3 times

Almylle 5 months, 3 weeks ago

Selected Answer: C

I labbed it and u can't configure af-interface in OSPFv3 address-family unicast routing, so the answer is C.

upvoted 2 times

yonig 8 months, 2 weeks ago

the correct answer is C.
answer D (even if there no sysntax error) does not adds interface, its just adds the family. the question states " adds a nother interface" meaning - the address family IPV4 unicast is already configured and the command to associate a new interface to OSFPv3 is in answer C

upvoted 1 times

Koume 11 months, 1 week ago

Selected Answer: C

The only method to add an interface on ospfv3 y by interface basis so C is correct

upvoted 1 times

nicoaburto 11 months, 2 weeks ago

D - because the configuration be applied into process OSPFv3 - D contain 2 commands

upvoted 1 times

wts 1 year, 3 months ago

Selected Answer: C

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/xe-3s/iro-xe-3s-book/ip6-route-ospfv3-add-fam-xe.html#:~:text=another%20routing%20domain.-,Enabling%20OSPFv3%20on%20an%20Interface,-SUMMARY%20STEPS

upvoted 1 times

Nhan 1 year, 3 months ago

Make it simple, the configuration is under an interface, not router configuration mode, A and D are not even relevant to the case

upvoted 2 times

Alexloh 1 year, 5 months ago

The answer is C, below the sample config for OSPFv3

```
R2(config)# router ospfv3 1
R2(config-router)# address-family ipv4 unicast
R2(config-router-af)# passive-interface Lo0
R2(config-router-af)# exit
R2(config-router)# exit
R2(config)# interface Loopback 0
R2(config-if)# ospfv3 1 ipv4 area 1
R2(config-if)# interface Serial 0/0
R2(config-if)# ospfv3 1 ipv4 area 1
```

upvoted 4 times

Iarn 1 year, 7 months ago

Selected Answer: C

C 100%

upvoted 1 times

xziomal9 1 year, 8 months ago

Selected Answer: C

The correct answer is: C

upvoted 1 times

The_KingPK 1 year, 9 months ago

C is Correct

upvoted 1 times

YaPet 1 year, 10 months ago

Selected Answer: C

C is correct.
From Cisco command reference examples:
Device(config-if)# ospfv3 1 area 1 ipv4 --- Enables OSPFv3 on an interface with the IPv4 or IPv6 AF.

address-family ipv4 unicast --- Enters IPv4 address family configuration mode for OSPFv3.

upvoted 1 times

JingleJangus 1 year, 10 months ago

Selected Answer: C

Correct answer is C

upvoted 1 times

Nonono 1 year, 10 months ago

C is correct

upvoted 1 times

Networkingguy 1 year, 10 months ago

Selected Answer: C

C is correct here

upvoted 1 times

Girmiti 1 year, 11 months ago

the correct answer should be C.

upvoted 1 times

Question #20 *Topic 1*

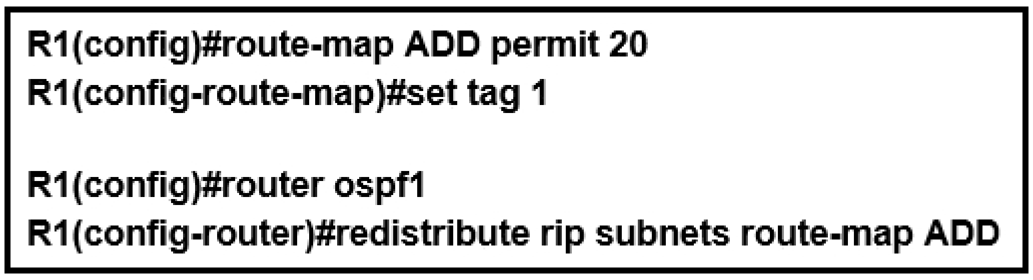

```
R1(config)#route-map ADD permit 20
R1(config-route-map)#set tag 1

R1(config)#router ospf1
R1(config-router)#redistribute rip subnets route-map ADD
```

Refer to the exhibit. Which statement about R1 is true?

A. OSPF redistributes RIP routes only if they have a tag of one.

B. RIP learned routes are distributed to OSPF with a tag value of one.

C. R1 adds one to the metric for RIP learned routes before redistributing to OSPF.

D. RIP routes are redistributed to OSPF without any changes.

Correct Answer: *B*

Community vote distribution

B (100%)

TigerDrev Highly Voted 3 years, 5 months ago

B is correct. If there is no match statement, it matches everything.

upvoted 9 times

Alexloh Most Recent 1 year, 5 months ago

Selected Answer: B

Agreed B is the correct answer.

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

ITBiscuit 2 years, 8 months ago

The answer is B --
https://www.cisco.com/c/en/us/td/docs/security/asa/asa84/configuration/guide/asa_84_cli_config/route_maps.pdf
"• If a match command or Match Clause value in ASDM is not present, all routes match the clause. In
the previous example, all routes that reach clause 30 match; therefore, the end of the route map is never reached."

upvoted 1 times

CraigB83 3 years, 2 months ago

If a match command is not present, all routes match the clause. In the previous example, all routes that reach clause 30 match; therefore, the end of the route-map is never reached.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/49111-route-map-bestp.html

upvoted 2 times

GustavoF 3 years, 5 months ago

B is the correct answer. Once there is no more configuration under the route-map and it's applied on the rip redistribution inside ospf, the router it is going to add a TAG 1 over in all route that came from RIP.

upvoted 3 times

heamgu 3 years, 6 months ago

In the exhibit, the route map route-map ADD permit 20 set tag 1... is not matching any ip, so the route map is not tagging the RIP routes when redistributed. Best answer for me is D.

upvoted 2 times

Question #21 Topic 1

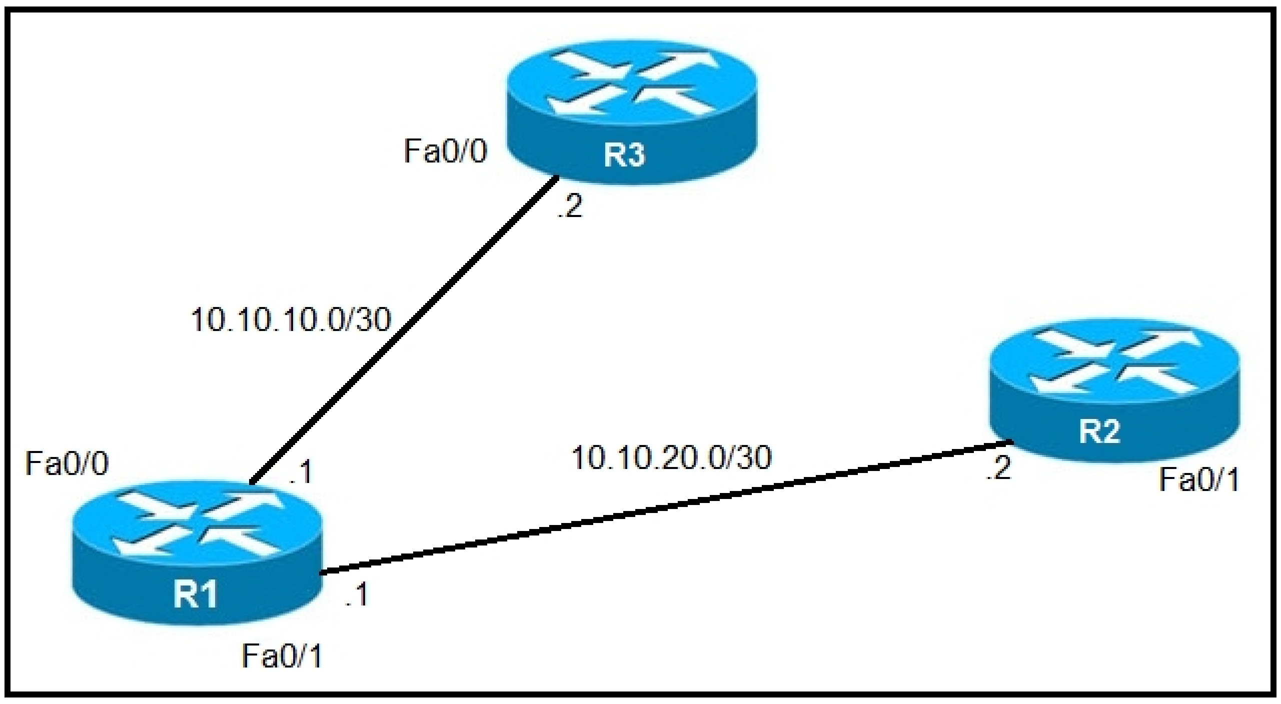

Refer to the exhibit. An IP SLA was configured on router R1 that allows the default route to be modified in the event that Fa0/0 loses reachability with the router R3

Fa0/0 interface. The route has changed to flow through router R2.

Which debug command is used to troubleshoot this issue?

A. debug ip flow

B. debug ip sla error

C. debug ip routing

D. debug ip packet

Correct Answer: *C*

Community vote distribution

C (73%) B (27%)

⊟ **jbr21** Highly Voted 2 years, 8 months ago

The answer is 'debug ip sla error' (C) -- The route has already changed, so debug IP routing is useless. We need to find out why the IP sla is failing and thus redirecting the default route to R2, as such we need to look at the current IP sla error debugging.

upvoted 9 times

⊟ **jbr21** 2 years, 8 months ago

B rather, not C -- whatever 'debug IP sla error' is the answer.

upvoted 4 times

⊟ **DonMike** Highly Voted 1 year, 10 months ago

C

The debug ip sla error command displays debug messages when an IP SLA run-time error occurs.
The debug ip sla error command can be used to troubleshoot problems that occur because of IP SLA misconfigurations or scheduler errors. Examples of problems that could cause IP SLA run-time errors include a disabled responder or a missing target.

upvoted 5 times

⊟ **tinoe** Most Recent 1 week ago

This question is incorrectly asked, otherwise it has no answer. Debug IP SLA ERROR is useless because the SLA does not have an error, it's working perfect by re-routing traffic to R2 (that is what it should be doing). Debug ip routing won't give any output if the change has already happened, so you cannot use it to troubleshoot the change that has already happened(it would have been useful if it was configured before the change). Debug ip packets gives no useful information and debug ip flow is just as useless as well.

upvoted 1 times

⊟ **diegodavid82** 4 months, 2 weeks ago

Selected Answer: B

debug ip sla error is the correct answer because debug ip routing is for troubleshooting routing protocols.

upvoted 2 times

⊟ **HungarianDish** 7 months ago

Selected Answer: C

These questions are often based on Cisco Press articles. If this is the relating article then answer "C" fits best.
https://www.ciscopress.com/articles/article.asp?p=1613547&seqNum=3

Scenario: Tracking Reachability to Two ISPs
Using "debug ip routing" for troubleshooting failed primary route. Output shows the route to be deleted, then missing.
upvoted 2 times

anonymous1966 8 months, 3 weeks ago

Selected Answer: C

Right answer C.
It cannot be B. Look at the output:
Router# debug ip sla error
May 5 05:00:35.483: control message failure:1
May 5 05:01:35.003: control message failure:1
May 5 05:02:34.527: control message failure:1
May 5 05:03:34.039: control message failure:1

Source: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/debug/command/i1/db-i1-cr-book/db-i3.html
upvoted 2 times

Pietjeplukgeluk 1 month, 1 week ago

If C was correct, would only make sense if the route was not already changed. "debug ip routing" only provides info, when a route is changing, and clearly the route has already be changed. Keep it at a bad question, all are wrong in a way. I personally stick with "debug ip sla error" as a correct answer, it creates shitty output, but that is better than nothing.
upvoted 1 times

Dominik_Networker 10 months ago

Selected Answer: B

B should be the correct answer
upvoted 1 times

Koume 11 months, 1 week ago

The best answer is "debug ip sla error" first they are talking about an IP sla that modifies the default route if theres is a fail on SLA, then stablished that traffic started flowing to R2 due to this config. This mean that there were an error on ip sla, so is failing and and a static floating route is installed, so you the core issue to verify why sla is failing and triggering the change. Using of "Debug ip routing" will now give any output as the route has already change.
upvoted 2 times

Dacusai 1 year, 4 months ago

They talking about an issue, so assume that R1 doesn't loose reachability to R3, the route change so you have to find out why
upvoted 1 times

timtgh 1 year, 6 months ago

If the route has changed to flow through router R2, then SLA is working. That's what it was supposed to do. So there is no SLA error. The error is whatever caused the unreachability that triggered the SLA to do its job.
upvoted 2 times

xziomal9 1 year, 8 months ago

Selected Answer: C

The correct answer is: C
upvoted 2 times

bayolo10 1 year, 8 months ago

Answer B
upvoted 2 times

Networkingguy 1 year, 10 months ago

Selected Answer: C

C looks to be the correct answer
upvoted 2 times

Networkingguy 1 year, 10 months ago

Sorry, Change this to B 'debug ip sla error'
upvoted 2 times

error_909 2 years, 3 months ago

The given answer is correct.

After testing GNS3:
The only result that make since is "debug ip routing"
upvoted 3 times

[Removed] 1 year, 11 months ago

How does it make sense if you're doing it AFTER the route has already been changed? Using debug ip sla error to verify the ip sla config has failed and then using debug ip sla trace afterwards are much better troubleshooting options...
upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

RHK0783 2 years, 8 months ago

Tricky :)

Question has referenced the IP SLA but asked to debug the routing which IP SLA command will only show the SLA triggering process. Route process can be looked at only using option C..

upvoted 4 times

thissiteisgreat 2 years, 11 months ago

Because the verification is hinged on whether the default route is replaced when the tracking default route is down, so debug ip routing is the right move.

upvoted 2 times

Question #22 *Topic 1*

Which configuration enables the VRF that is labeled `Inet` on FastEthernet0/0?

A. R1(config)# ip vrf Inet R1(config-vrf)#ip vrf FastEthernet0/0

B. R1(config)#ip vrf Inet FastEthernet0/0

C. R1(config)# ip vrf Inet R1(config-vrf)#interface FastEthernet0/0 R1(config-if)#ip vrf forwarding Inet

D. R1(config)#router ospf 1 vrf Inet R1(config-router)#ip vrf forwarding FastEthernet0/0

Correct Answer: *C*

Community vote distribution

C (100%)

Alexloh 1 year, 5 months ago

Selected Answer: C

C is the correct answer

upvoted 2 times

xziomal9 1 year, 8 months ago

Selected Answer: C

The correct answer is: C

upvoted 1 times

Girmiti 1 year, 11 months ago

Selected Answer: C

C is correct if R1(config-vrf)#interface FastEthernet0/0 will excluded.

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

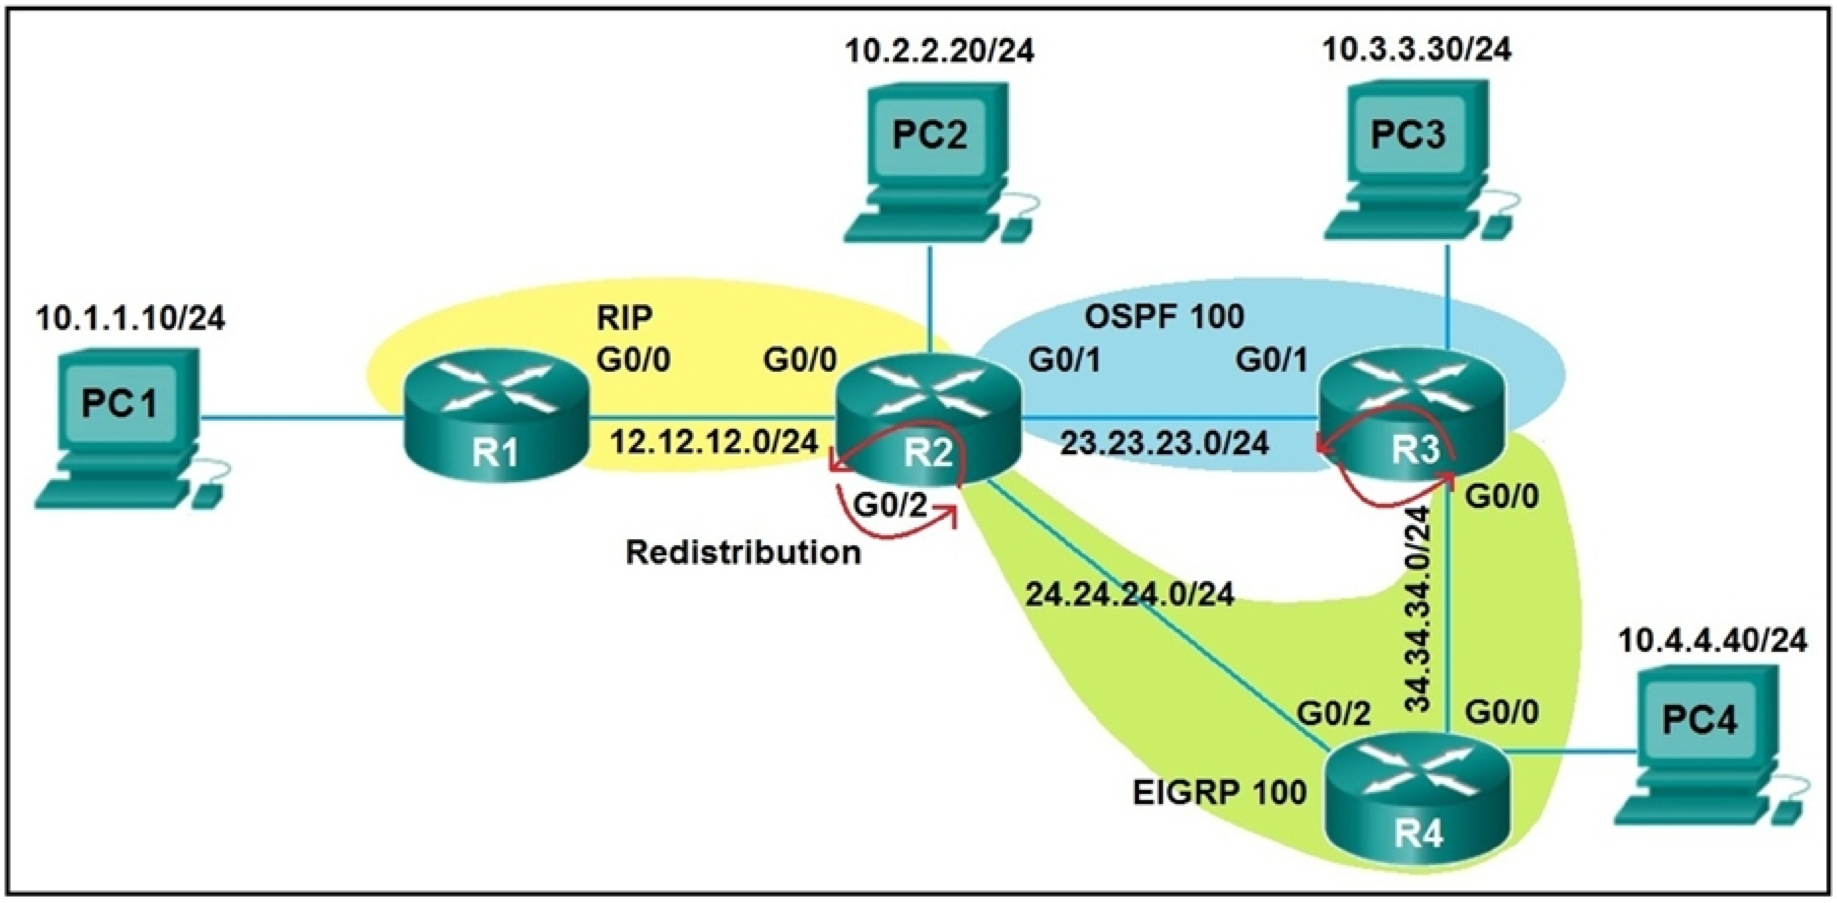

Refer to the exhibit. After redistribution is enabled between the routing protocols; PC2, PC3, and PC4 cannot reach PC1.
Which action can the engineer take to solve the issue so that all the PCs are reachable?

A. Set the administrative distance 100 under the RIP process on R2.

B. Filter the prefix 10.1.1.0/24 when redistributed from OSPF to EIGRP.

C. Filter the prefix 10.1.1.0/24 when redistributed from RIP to EIGRP.

D. Redistribute the directly connected interfaces on R2.

Correct Answer: *B*

Community vote distribution

A (94%) 6%

HETKAR Highly Voted 2 years, 9 months ago

This Config works: Answer A

```
-------------------------------------
R2#sh run | s rip
redistribute rip metric 1 1 1 1 1
router rip
version 2
redistribute eigrp 100 metric 1
network 10.0.0.0
network 12.0.0.0
distance 100
no auto-summary
-----------------------------------
R3#sh run | s router
router eigrp 100
network 34.34.34.0 0.0.0.255
redistribute ospf 100 metric 1 1 1 1 1
router ospf 100
redistribute eigrp 100 subnets
network 10.3.3.0 0.0.0.255 area 0
network 23.23.23.0 0.0.0.255 area 0
---------------------------------------
```

Answer B is wrong: the Correct is to filter 10.1.1.10 when redistribute from EIGRP to OSPF: Configs are

```
---------------------------------------
ip prefix-list DNA seq 5 deny 10.1.1.0/24
ip prefix-list DNA seq 10 permit 0.0.0.0/0 le 32
route-map DDD permit 10
match ip address prefix-list DNA
!
router eigrp 100
network 34.34.34.0 0.0.0.255
redistribute ospf 100 metric 1 1 1 1 1
!
```

```
router ospf 100
redistribute eigrp 100 subnets route-map DDD
network 10.3.3.0 0.0.0.255 area 0
network 23.23.23.0 0.0.0.255 area 0
!
```

upvoted 20 times

Alnet 2 years ago

100% agree. Labbed it. Reducing AD to 100 will always provide an exit for packets to 10.1.1.0/24. When this route is in EIGRP it's external, so it will be treated with AD=170. Within OSPF AD=110. Reduce RIP down to 100 then on R2 10.1.1.0/24 will always point out towards RIP domain.
After making lab you'll see that problem happens on R2; it redistributes RIP into EIGRP, which then gets redistributed into OSPF at R3. So R2 learns from R3 an OSPF E2 route with an AD of 110. It inserts 10.1.1.0/24 >> R3 into the RIB because the OSPF AD is lower than the RIP learned AD.
Thus lower RIP AD to 100 and it will be preferred over the OSPF route.

upvoted 10 times

myrmike 1 year, 11 months ago

What is being redistributed on R2? I may be missing something but when I labbed the below all routers could ping the 10.1.1.1 interface on R1. There were no redistributions on R4.

```
R2(config)#do sho run | s router
router eigrp 100
network 24.24.24.2 0.0.0.0
redistribute rip metric 1000000 10 255 1 1500
router ospf 100
redistribute rip
network 23.23.23.2 0.0.0.0 area 0
router rip
version 2
redistribute ospf 100 metric 2
redistribute eigrp 100 metric 2
network 10.0.0.0
network 12.0.0.0
neighbor 12.12.12.1
R2(config)#


R3#sho run | s router
router eigrp 100
network 34.34.34.0 0.0.0.255
redistribute ospf 100 metric 1000000 10 255 1 1500
router ospf 100
redistribute eigrp 100
network 10.3.3.3 0.0.0.0 area 0
network 23.23.23.3 0.0.0.0 area
```

upvoted 1 times

kent2612 1 year, 10 months ago

Me too I lab it up and there's no issue. PC2, PC3 & PC4 could ping PC1

```
R2#show run | s router
router eigrp 100
redistribute rip metric 1000000 1 255 1 1500
redistribute ospf 100 metric 1000000 1 255 1 1500
network 24.24.24.0 0.0.0.255
no auto-summary
router ospf 100
log-adjacency-changes
redistribute rip subnets
redistribute eigrp 100 subnets
network 23.23.23.0 0.0.0.255 area 0
router rip
version 2
redistribute ospf 100 metric 1
redistribute eigrp 100 metric 1
network 10.0.0.0
network 12.0.0.0
no auto-summary

R3#show run | s router
router eigrp 100
redistribute ospf 100 metric 1000000 1 255 1 1500
network 34.34.34.0 0.0.0.255
no auto-summary
router ospf 100
log-adjacency-changes
redistribute eigrp 100 subnets
network 10.3.3.0 0.0.0.255 area 0
network 23.23.23.0 0.0.0.255 area 0
```

upvoted 2 times

quyle 1 year, 2 months ago

I lab all router can ping PC1 =)))), maybe question is not true

upvoted 1 times

larn 1 year, 7 months ago

You still will have a routing loop

upvoted 1 times

larn 1 year, 7 months ago

PC 2 will not be able to reach PC1 The correct answer is B

upvoted 1 times

uglyprawn 2 years, 9 months ago

very good. i dont need to test this one to understand. answer is A

upvoted 5 times

HieuPham 2 years, 5 months ago

What's result your config?
B correct: It seems there is a loop because of mutual redistributions among RIP, OSPF and EIGRP domains. So we should filter out the prefix 10.1.1.0/24 when redistributed from OSPF to EIGRP (the second redistribution point) to prevent routing loop.

upvoted 3 times

Jenia1 Highly Voted 1 year, 10 months ago

Selected Answer: A

HETKAR and Alnet are correct, and I just want to add simple clarification
D - does not make any sense
C - if you filter prefix 10.1.1.0/24 from RIP to EIGRP, this network becomes unreachable on R3 and R4.
B - there is no redistribution from RIP to OSPF that won't work as OSPF does not learn the prefix from RIP.
A (Correct) When the traffic goes to the R2, the router will have a choice - sent to the R1 (RIP AD is 120) or to R3 (OSPF 110). OSPF is learning the route from R4 via redistribution.
Route with the lover AD will be injected into the routing table.
So if RIP's AD will not be changed to 100, R2 will forward the traffic to R3, so the packet will not reach R1.
I was confused a bit when I saw this scheme first time. R2 OSPF is redistributed into the RIP and RIP redistributed into EIGRP
I hope it helps

upvoted 7 times

Wooker Most Recent 9 months ago

Selected Answer: A

Answer A

upvoted 1 times

Koume 11 months, 1 week ago

Selected Answer: A

The core issue here is that When the RIP route redistributed from eigrp into OSPF on R3, R2 that is running OSPF will install the route as have better AD, causing the loop. The solution here then is use a distribute list to do not intall the EIGRP route ont the ospf procees and avoid the loop.

upvoted 1 times

ChillingAgain 1 year, 1 month ago

Selected Answer: A

Redistribution of subnet 10.1.1.0/24 on from RIP to OSPF on R2 with create an OSPF route to 10.1.1.0/24 with AD of 110. This one is preffered over the RIP route to 10.1.1.0/24 with AD 120.
Redistribution of RIP route 10.1.1.0/24 to EIGRP on R2 creates an external EIGRP route with AD 170. This route will not be chosen anyway.
So if you set the AD of RIP to 100 on R2 that route is chosen to reach 10.1.1.0/24.

upvoted 3 times

Edwinmolinab 1 year, 1 month ago

Selected Answer: B

Given answer is correct. I was testing on GNS3 and is the best solution

upvoted 1 times

wts 1 year, 3 months ago

It may seem that B solves the problem, which makes it difficult to choose an answer.

But B removes only the EIGRP route that passed from R2 in a clockwise direction.
The main problem is that on R2 there is an external OSPF route(counterclock-wise) that pulls all attempts to get to the PC1.

If in answer B we swapped OSPF and EIGRP, then it would fit. .

upvoted 1 times

WAKIDI 1 year, 5 months ago

if the red arrow in the picture is a symbol to a redistribution, what we should have in R2 are : OSPF is redistributed into RIP, RIP is redistributed into EIGRP and an arrow between OSPF and EIGRP that i can't see where it is pointing at.

upvoted 1 times

timtgh 1 year, 6 months ago

A - solves the problem because R2 trusts RI and sends the 10.1.1.0 traffic to the left.
B - doesn't help because the problem is caused by redistributing the route from EIGRP to OSPF, not the other way around.

C - suppresses the 10.1.1.0 route from EIGRP, thereby preventing PC4 from reaching the subnet.
D - just ridiculous nonsense obviously.

upvoted 3 times

timtgh 1 year, 6 months ago

typo - first line should say R2 trusts RIP

upvoted 1 times

larn 1 year, 7 months ago

This is a route looping problem, being the route is looping via redistribution from OPSF to EIGRP. Why would adding distance metric to RIP solve this?

upvoted 1 times

timtgh 1 year, 6 months ago

Because when R2 is trying to get to 10.1.1.1, it will always go LEFT to the correct destination if it trusts the RIP routes over the OSPF routes. With RIP having AD of 100 it is trusted over OSPF which is 110.

upvoted 1 times

xziomal9 1 year, 8 months ago

Selected Answer: A

The correct answer is: A

upvoted 1 times

DonMike 1 year, 10 months ago

Looks like no answer is correct. Just labbed it. It works when you set an AD of 171 for OSPF routes on R2. But once the RIP distance is set to 100 on R2 instead R1 loses connectivity to PC3 because there is no redistribution from OSPF into RIP and the best route to PC3 in R2s routing table is via OSPF. So all routes from R3 must traverse EIGRP (since this is redistributed into RIP afterwards) which means that these routes must be in R2s routing table available via EIGRP. Nonetheless A makes most sense.

```
r2#show running-config | s router
router eigrp 1
default-metric 1000000 1 255 1 1500
network 24.24.24.0 0.0.0.255
redistribute rip route-map set-rip-tag
router ospf 1
network 23.23.23.0 0.0.0.255 area 0
distance 171
router rip
version 2
redistribute eigrp 1 metric 1
network 10.0.0.0
network 12.0.0.0
no auto-summary
```

upvoted 1 times

JingleJangus 1 year, 11 months ago

Selected Answer: A

Definitely A.

upvoted 2 times

testbench007 1 year, 11 months ago

A is the correct answer. OSPF has the AD of 110 and RIP has the AD of 120. R2 will select the OSPF learnt info and eject the RIP route. Once the RIP route has been ejected then EIGRP cannot redistribute the 10.1.1.0/24 into EIGRP since RIP no longer has it.

upvoted 1 times

geek1992 1 year, 11 months ago

A i just tested

upvoted 1 times

CiscoSystems 2 years, 1 month ago

B is correct as 10.1.1.0/24 is redistributed into both OSPF and EIGRP at R2. This causes a routing loop as R3 is redistributing between OSPF and EIGRP. So to prevent the routing loop, you can filter the prefix 10.1.1.0/24 from OSPF to EIGRP. Changing the AD to 100 for RIP on R2 won't affect the route 10.1.1.0/24 learned on R3 and R4, because it's learned from OSPF/EIGRP as it's redistributed. Not learned directly from RIP. It will also not matter on R2 because the only route installed in the RIB will be from RIB, assuming that the routing loop is prevented with the aforementioned prefix-list.

upvoted 1 times

Alnet 2 years, 1 month ago

Where is the loop? If it's between R3 R4 (bouncing back and forth), then you need to filter it to prevent, hands down, no question. But if the path keeps going around the triangle, then the answer is to give it a way to get out of the triangle. To get it to exit the triangle you can lower the RIP AD to 100 (100, 110, 170 > 100 wins) and the RIP route then gets installed in the RIB at R2 providing a way out of the triangle. But again, if your traffic bounces back and forth between R3 > R4, then making changes on R2 will do nothing for you. Also changing AD keeps multiple paths in everyones tables, filtering makes sure you only have one way to get there (no failover path).

upvoted 1 times

kent2612 1 year, 9 months ago

I don't see changing RIP AD to 100 helps as R2 will still route traffic via R1 even when I set the AD to 100 or 180.

```
R2#sh ip route | in 10.1.1.0
```

```
R 10.1.1.0 [120/1] via 12.12.12.1, 00:00:38, GigabitEthernet0/0
R2#show ip eigrp topology 10.1.1.0/24
IP-EIGRP (AS 100): Topology entry for 10.1.1.0/24
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2560256
Routing Descriptor Blocks:
12.12.12.1, from Redistributed, Send flag is 0x0
Composite metric is (2560256/0), Route is External
Vector metric:
Minimum bandwidth is 1000 Kbit
Total delay is 10 microseconds
Reliability is 255/255
Load is 1/255
Minimum MTU is 1500
Hop count is 0
External data:
Originating router is 24.24.24.2 (this system)
AS number of route is 0
External protocol is RIP, external metric is 1
Administrator tag is 0 (0x00000000)
```

upvoted 1 times

⊟ **CiscoSystems** 2 years, 1 month ago

B is correct as 10.1.1.0/24 is redistributed into both OSPF and EIGRP at R2. This causes a routing loop as R3 is redistributing between OSPF and EIGRP. So to prevent the routing loop, you can filter the prefix 10.1.1.0/24 from OSPF to EIGRP at R3.

upvoted 2 times

Question #24 Topic 1

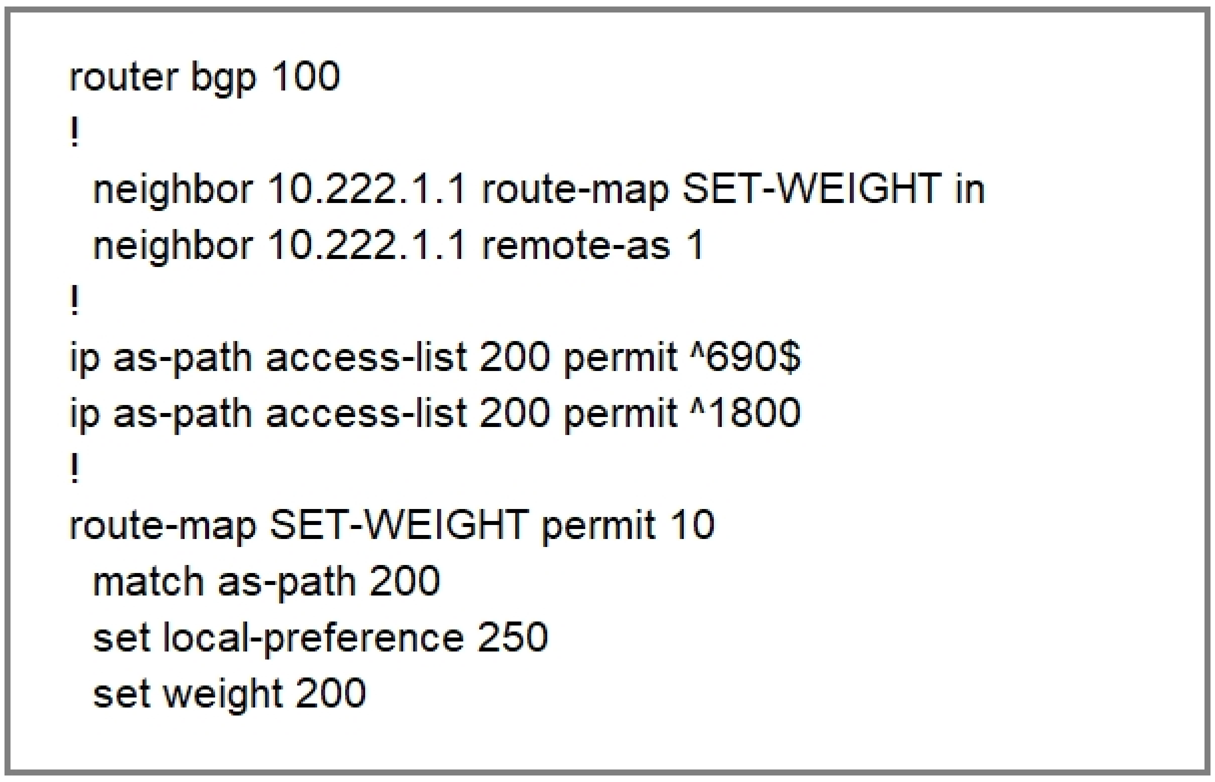

```
router bgp 100
!
  neighbor 10.222.1.1 route-map SET-WEIGHT in
  neighbor 10.222.1.1 remote-as 1
!
ip as-path access-list 200 permit ^690$
ip as-path access-list 200 permit ^1800
!
route-map SET-WEIGHT permit 10
  match as-path 200
  set local-preference 250
  set weight 200
```

Refer to the exhibit. A router is receiving BGP routing updates from multiple neighbors for routes in AS 690.
What is the reason that the router still sends traffic that is destined to AS 690 to a neighbor other than 10.222.1.1?

A. The local preference value in another neighbor statement is higher than 250.

B. The local preference value should be set to the same value as the weight in the route map.

C. The route map is applied in the wrong direction.

D. The weight value in another neighbor statement is higher than 200.

Correct Answer: *D*

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3se/3850/irg-xe-3se-3850-book/irg-prefix-filter.html

Community vote distribution

D (100%)

⊟ **JingleJangus** Highly Voted 1 year, 11 months ago

We=weight
love=local preference
oranges= originated locally >remotely
as= as path
oranges= origin code (IGP, EGP, incomplete?)
mean= MED (metric)
pure= eBGP > iBGP learned
refreshment= highest RID.

upvoted 11 times

⊟ **Earl03** Highly Voted 3 years, 5 months ago

Should be a, as weight values aren't shared between routers, but local preference is, correct?

upvoted 5 times

⊟ **geek1992** 2 years, 12 months ago

Local Pref is shared in the same AS

upvoted 4 times

⊟ **timtgh** 1 year, 6 months ago

Sharing is irrelevant. The question is about what THIS router does. We don't know (or care) if there even are other routers in this AS. On THIS router, we have assigned a weight to one neighbor, and (according to the correct answer) we have assigned a higher weight to a different neighbor.

upvoted 2 times

⊟ **Ll123123** Most Recent 2 months ago

The AS of the router is 100, the neighbor 10.22.22.1 is AS 1, then the route announce to this must be eBGP route, and it must have a 1 prepend in it. The patten ^690$ shall match any path list that start with 690 and end with 690, so it should not set any weight nor local preference on that.. while local preference does not have effect, the weight should make it prefer to other, but it should not be advertised by other but rather have another setting to set the WEIGHT higher..

upvoted 1 times

⊟ **MD_Shox** 1 year ago

show ip bgp regexp ^100$ match the direct peering learned routes from as 100
^$ - match locally originated routes
only D makes sense

upvoted 1 times

⊟ **wts** 1 year, 3 months ago

Not only is the situation described disgustingly, but also regexp will not match prefixes from AS1.
We do not see the settings of "another neighbors".

The question is obviously on the knowledge of path selection, but it is not clear where to apply this knowledge here ...

upvoted 1 times

⊟ **xziomal9** 1 year, 8 months ago

Selected Answer: D

The correct answer is: D

upvoted 1 times

⊟ **wts** 1 year, 9 months ago

Why does the neighbor have AS1, and we are waiting for updates from AS690?

upvoted 1 times

⊟ **lcy1** 1 year, 10 months ago

correct answer is not in options - router is sending traffic elsewhere, because it can never receive update matching as-path statement ^690$ from neighbor in AS 1. So it must receive update from other neighbors and that's why it sends it the other way. But when following "excluding wrong options" approach, D remains the last.

upvoted 3 times

⊟ **YaPet** 1 year, 10 months ago

Just D seems to be correct, because LOCAL-PREFERENCE is used for choosing best path between different routers.

upvoted 1 times

⊟ **Networkingguy** 1 year, 10 months ago

Selected Answer: D

D is correct because we are only dealing with one router here, with its different bgp route statements.

upvoted 1 times

⊟ **Stivostine** 2 years ago

Attributes are processed in the order :

1. Prefer the highest weight
2. Prefer the highest local preference

D is ok

upvoted 1 times

⊟ **JOKERR** 2 years ago

I tested in GNS3 and router is choosing weight first between 2 eBGP neighbors for the same route.

upvoted 1 times

⊟ **gndrx78** 2 years ago

Selected Answer: D

weight comes before local-pref in BGP routing decisional process

upvoted 3 times

⊟ **CiscoSystems** 2 years, 1 month ago

D is correct

upvoted 1 times

⊟ **Alnet** 2 years, 1 month ago

We all know Weight is local to each router, right? So if one of the neighbors has already been set to a weight of 200 (which one has in this config), then the only reason a third neighbor would be used is if weight of a third neighbor was higher.
Because weight is more preferred than Local Pref, it won't matter what you set Local Pref to on any neighbor, the chosen one will always be the weight=200 UNLESS someone else has a higher weight.
And don't forget, we're only looking at this one router, not a set of routers in an AS.

upvoted 2 times

⊟ **kuzma** 2 years, 1 month ago

ip as-path access-list 200 permit ^690$ - prefixes originated in our neighbor AS 690
but neighbor 10.222.1.1 remote-as 1
This route-map will no be in use.

upvoted 4 times

⊟ **AliMo123** 2 years, 1 month ago

it is a wrong question as you stated above
AS 1 and then AS 690 do not make any sense here

upvoted 1 times

donjime 2 years, 2 months ago

D is the correct Answer

upvoted 1 times

Question #25 *Topic 1*

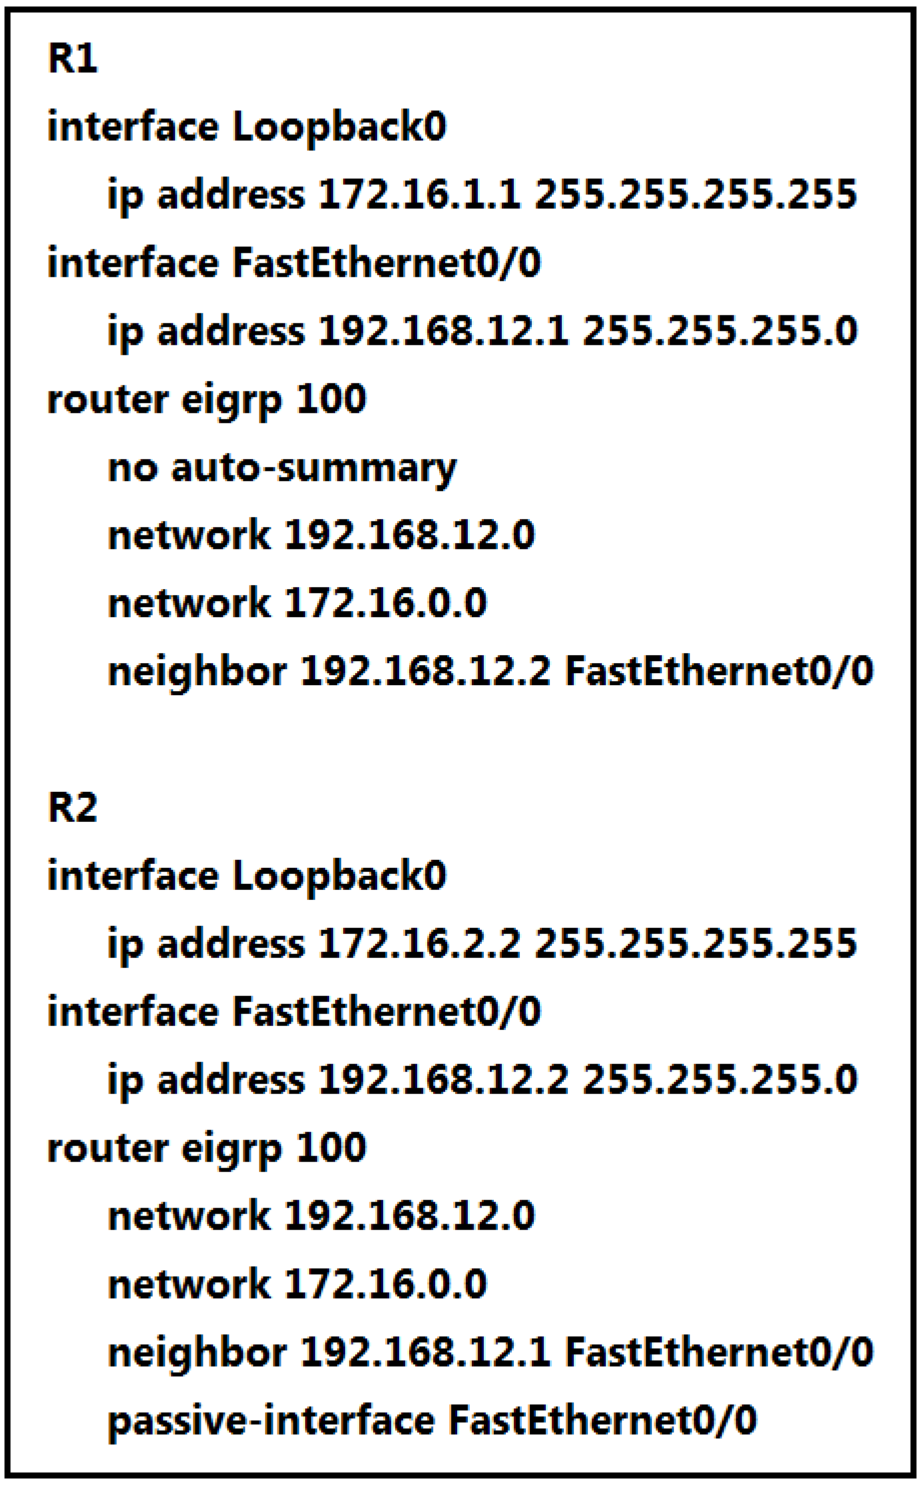

```
R1
interface Loopback0
    ip address 172.16.1.1 255.255.255.255
interface FastEthernet0/0
    ip address 192.168.12.1 255.255.255.0
router eigrp 100
    no auto-summary
    network 192.168.12.0
    network 172.16.0.0
    neighbor 192.168.12.2 FastEthernet0/0

R2
interface Loopback0
    ip address 172.16.2.2 255.255.255.255
interface FastEthernet0/0
    ip address 192.168.12.2 255.255.255.0
router eigrp 100
    network 192.168.12.0
    network 172.16.0.0
    neighbor 192.168.12.1 FastEthernet0/0
    passive-interface FastEthernet0/0
```

Refer to the exhibit. R1 and R2 cannot establish an EIGRP adjacency.
Which action establishes EIGRP adjacency?

A. Remove the current autonomous system number on one of the routers and change to a different value.

B. Add the passive-interface command to the R1 configuration so that it matches the R2 configuration.

C. Remove the passive-interface command from the R2 configuration so that it matches the R1 configuration.

D. Add the no auto-summary command to the R2 configuration so that it matches the R1 configuration.

Correct Answer: *C*

Community vote distribution

C (100%)

⊟ **Alexloh** 1 year, 5 months ago

Selected Answer: C

The correct answer is C

upvoted 1 times

⊟ **xziomal9** 1 year, 8 months ago

Selected Answer: C

The correct answer is: C

upvoted 1 times

⊟ **studybuddy10** 2 years, 1 month ago

Given answer is correct, labbed and as soon as you configure passive-interface it tears down the neighbor as below.

EIGRP-IPv4 100: Neighbor 192.168.12.1 (Ethernet0/0) is down: interface passive

upvoted 4 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #26 *Topic 1*

An engineer configured policy-based routing for a destination IP address that does not exist in the routing table.
How is the packet treated through the policy for configuring the set ip default next-hop command?

A. Packets are not forwarded to the specific next hop.

B. Packets are forwarded based on the routing table.

C. Packets are forwarded based on a static route.

D. Packets are forwarded to the specific next hop.

Correct Answer: *A*

Community vote distribution

D (90%) 5%

Alex147 Highly Voted 👍 2 years ago

Selected Answer: D

D

The set ip default next-hop command verifies the existence of the destination IP address in the routing table, and...

if the destination IP address exists, the command does not policy route the packet, but forwards the packet based on the routing table.

if the destination IP address does not exist, the command policy routes the packet by sending it to the specified next hop.

The set ip next-hop command verifies the existence of the next hop specified, and...

if the next hop exists in the routing table, then the command policy routes the packet to the next hop.

if the next hop does not exist in the routing table, the command uses the normal routing table to forward the packet.

upvoted 9 times

SnoopDD Most Recent 2 months ago

A is correct

upvoted 1 times

HungarianDish 7 months ago

Selected Answer: D

when the destination route is not in the routing table, the packet is policy routed (to the specified next hop)
https://www.cisco.com/c/en/us/support/docs/ip/ip-routed-protocols/47121-pbr-cmds-ce.html#anc12

upvoted 2 times

HungarianDish 6 months, 3 weeks ago

set ip default next-hop:
-if destination IP not in RIB -> policy route
set ip next-hop:
-if destination IP not in RIB -> use normal routing table

upvoted 2 times

AinsB 7 months ago

Selected Answer: A

At first glance D would seem to be correct but if you think about it, to get to a path it must be known and remember there is the RIB and FIB. So if it is not known then it is not in the RIB and the default action is drop or send to the Default Gateway

upvoted 1 times

anonymous1966 8 months, 3 weeks ago

Selected Answer: B

In my opinion is "B".
This document provides a sample configuration for policy-based routing (PBR) with the set ip default next-hop and set ip next-hop commands.

The set ip default next-hop command verifies the existence of the destination IP address in the routing table, and:

if the destination IP address exists, the command does not policy route the packet, but forwards the packet based on the routing table.

if the destination IP address does not exist, the command policy routes the packet and sends it to the specified next hop.

The set ip next-hop command verifies the existence of the next hop specified, and:

if the next hop exists in the routing table, then the command policy routes the packet to the next hop.

if the next hop does not exist in the routing table, the command uses the normal routing table to forward the packet.
Source: https://www.cisco.com/c/en/us/support/docs/ip/ip-routed-protocols/47121-pbr-cmds-ce.html

upvoted 1 times

Stylar 9 months ago

Selected Answer: D

ChatGPT: Yes, if the router has been configured with a policy-based routing (PBR) rule using the "set ip default next-hop" command and a packet arrives at the router with a destination IP address that is not present in the router's routing information base (RIB), the router will forward the packet to the next-hop address specified in the PBR rule.

This is because PBR allows the router to apply forwarding policies that are independent of the routing table lookup process. In other words, the router will use the PBR policy to determine where to forward the packet, regardless of whether the destination IP address is present in the RIB or not.

However, it's important to note that forwarding packets using PBR rules that reference non-existent destinations can result in unexpected behavior and can lead to packet loss if the next-hop address specified in the PBR rule is not reachable. It's generally recommended to ensure that all destination IP addresses referenced in PBR rules are present in the RIB to avoid any unexpected packet drops.

upvoted 3 times

Noproblem22 1 year, 1 month ago

D is corrected, with "set ip default next-hop x.x.x.x" if there is no specific route on the routing table, it will use PBR.

upvoted 1 times

tipama7298 1 year, 2 months ago

D. The set ip default next-hop command verifies the existence of the destination IP address in the routing table, and...
if the destination IP address exists, the command does not policy route the packet, but forwards the packet based on the routing table.
if the destination IP address does not exist, the command policy routes the packet by sending it to the specified next hop.

The set ip next-hop command verifies the existence of the next hop specified, and...

if the next hop exists in the routing table, then the command policy routes the packet to the next hop.

if the next hop does not exist in the routing table, the command uses the normal routing table to forward the packet.

upvoted 2 times

Router 1 year, 3 months ago

d is the ans, policy base routing overrides the routing table

upvoted 1 times

Edwinmolinab 1 year, 5 months ago

Selected Answer: D

I Tested it on GNS3 and the packet was forwarded to the specific next hop, and the route wasn't in the routing table and not default gateway for the network

upvoted 1 times

larn 1 year, 7 months ago

Selected Answer: D

The destination IP/Subnet is not in the routing table, NOT the next hop IP address.

upvoted 1 times

xziomal9 1 year, 8 months ago

Selected Answer: D

The correct answer is: D

upvoted 1 times

YaPet 1 year, 10 months ago

Selected Answer: D

D is correct

upvoted 2 times

Jenia1 1 year, 10 months ago

Almost everyone chooses D, can you please clarify?
The statement is "An engineer configured policy-based routing for a destination IP address that does not exist in the routing table."
It's mean that RIB contains the reachable route and PBR is configured to the next hop that does not exist, so the traffic should not be forwarded to the next hop that does not exist, especially with "ip default next-hop command" as it checks the RIB first.
D. Packets are forwarded to the specific next hop.
According to the question, A should be the correct one, please let me know if I misunderstood the question, thanks

upvoted 2 times

bogd 1 year, 10 months ago

"destination IP address does NOT exist in the routing table" == "RIB does NOT contain the reachable route". Hence, the default next hop is used.

upvoted 1 times

timtgh 1 year, 6 months ago

Question is worded poorly. They mean that the destination is not in the routing table. There is nothing wrong with the configured next hop. The "set default next hop" feature means use the configured next hop only when the destination address in the packet is absent from the routing table, which is what the question says is happening.

upvoted 1 times

Carl1999 1 year, 10 months ago

The set ip default next-hop command verifies the existence of the destination IP address in the routing table.
https://www.cisco.com/c/en/us/support/docs/ip/ip-routed-protocols/47121-pbr-cmds-ce.html

upvoted 1 times

Carl1999 1 year, 10 months ago

This document provides a sample configuration for policy-based routing (PBR) using the set ip default next-hop and set ip next-hop commands.

The set ip default next-hop command verifies the existence of the destination IP address in the routing table, and...

if the destination IP address exists, the command does not policy route the packet, but forwards the packet based on the routing table.

if the destination IP address does not exist, the command policy routes the packet by sending it to the specified next hop.

upvoted 2 times

[Removed] 1 year, 10 months ago

It will still send it out to the specified next hop unless you add verify-availability to the command.

upvoted 1 times

Stivostine 2 years ago

D because :

If and only if there is no specific route in the routing table. Read that sentence
again. Why does this happen if and only if there is no specific route in the routing table?
Because the ip default next-hop command is used, PBR examines the routing table, and if
there is a specific route in the routing table, it is used. If there is no specific route in the routing table, the packet is routed using policy-based routing.

upvoted 2 times

studybuddy10 2 years, 1 month ago

D. Best simple explanation i have found is here:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/xe-3s/iri-xe-3s-book/iri-pbr-default-nexthop-route.pdf

From page 3.
"The set ip next-hop and set ip default next-hop commands are similar but have a different order of operation. Configuring the set ip next-hop command causes the system to first use policy routing and then use the routing table. Configuring the set ip default next-hop command causes the system to first use the routing table and then the policy-route-specified next hop"

In our case the route for the destination is not in the routing table therefor the configured next hop will be used.

upvoted 4 times

gndrx78 2 years ago

something similar is in "CCNP Enterprise Advanced Routing ENARSI 300-410 Official Cert Guide by Raymond Lacoste, Brad Edgeworth", Trouble Ticket 15-1

upvoted 1 times

AliMo123 2 years, 1 month ago

D is 100% correct since PBR overwrites RIB

upvoted 1 times

[Removed] 2 years, 1 month ago

You could change the wording and then B would be correct... set ip default next-hop - the router will check the routing table first and if no route then forwards to the specified next hop address. Thats why its D...

upvoted 1 times

Question #27 Topic 1

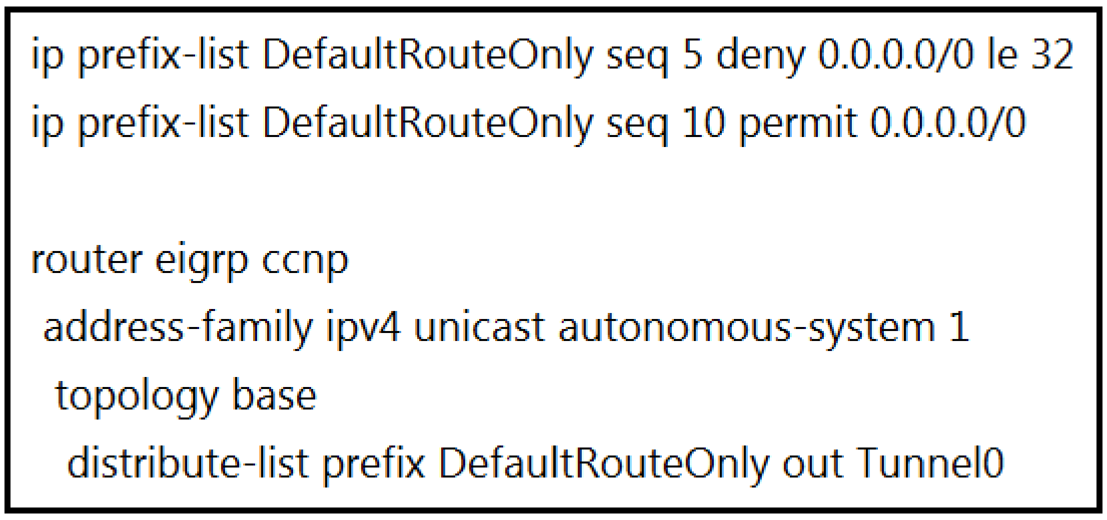

```
ip prefix-list DefaultRouteOnly seq 5 deny 0.0.0.0/0 le 32
ip prefix-list DefaultRouteOnly seq 10 permit 0.0.0.0/0

router eigrp ccnp
 address-family ipv4 unicast autonomous-system 1
  topology base
   distribute-list prefix DefaultRouteOnly out Tunnel0
```

Refer to the exhibit. The administrator configured route advertisement to a remote low resources router to use only the default route to reach any network but failed.
Which action resolves this issue?

A. Remove the prefix keyword from the distribute-list command.

B. Remove the line with the sequence number 10 from the prefix list.

C. Change the direction of the distribute-list command from out to in.

D. Remove the line with the sequence number 5 from the prefix list.

Correct Answer: *D*

Community vote distribution

D (100%)

Ll123123 1 month, 3 weeks ago

Selected Answer: D

deny 0.0.0.0/0 le 32 in prefix list means "any" route
permit 0.0.0.0/0 in prefix list means default route
So remove seq 5 will means only permit a default route advertise to the remote-host via the tunnel.
I don't get the meaning of the question because don't know if the config is the "remote-host". A picture can help better explain this question
upvoted 1 times

HungarianDish 8 months ago

https://community.cisco.com/t5/routing/very-quick-question-on-prefix-list-0-0-0-0-0/td-p/1356083
ip prefix-list 0.0.0.0/0 just matches the default-route not all routes. So that prefix-list filters out all routes except the default-route.
upvoted 2 times

WAKIDI 1 year, 5 months ago

Can Anyone explain what is this "remote low resources router" trying to do ? what does "to use only the default route to reach any network" really mean ?
upvoted 1 times

David98898998 6 months, 3 weeks ago

It means to keep the routing table small by only advertising to it a single default route. It will use this route to send all traffic.
upvoted 2 times

Nhan 1 year, 6 months ago

Basically deny/32 mean deny all route becuase ipv4 is 32 bit
upvoted 1 times

timtgh 1 year, 6 months ago

All answers are wrong. The permit statement permits all routing updates because it has a /0 mask. If they wanted to permit only default routes, the syntax should be 0.0.0.0/32, meaning all 3 bits of the advertised route must match 0.0.0.0. Also they should add le 0 at the end.
upvoted 1 times

timtgh 1 year, 6 months ago

Typo - "all 3 3bits" should say all 32 bits.
upvoted 1 times

luisdzrz 1 year, 6 months ago

El rango asumido para ge y le si no se especifica nada es 32
upvoted 1 times

timtgh 1 year, 6 months ago

Please disregard comment, was mixed up. /0 is correct.

upvoted 1 times

niveaking 1 year, 10 months ago

Selected Answer: D

Answer D is correct

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 3 times

Question #28

Topic 1

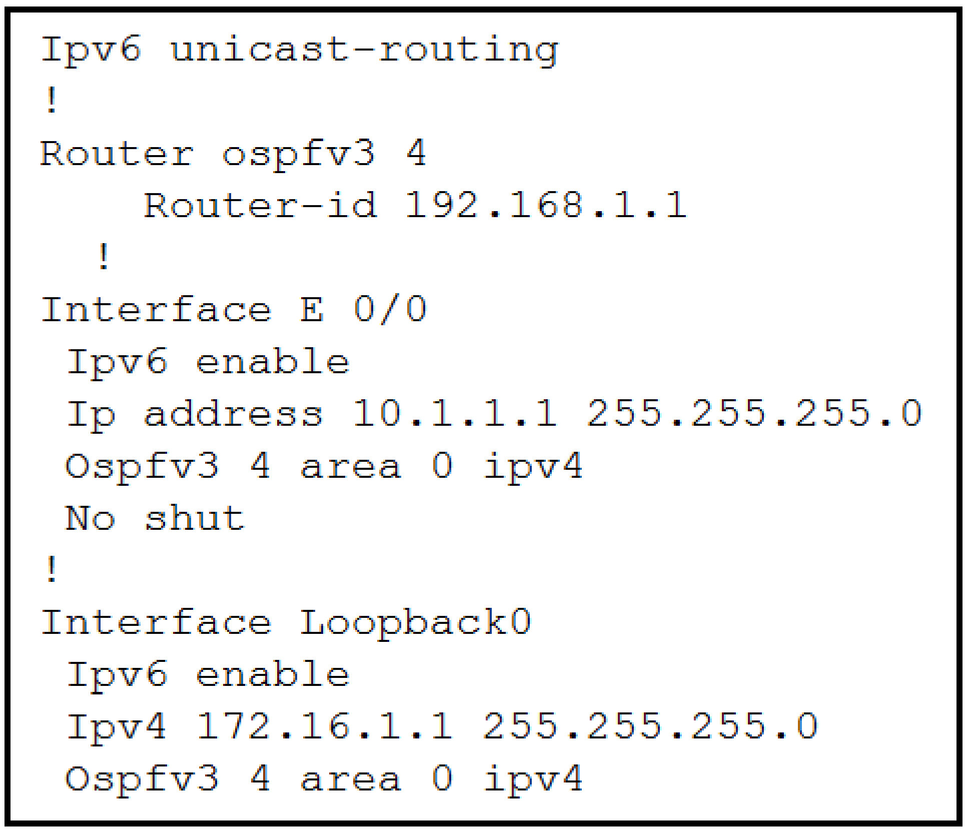

```
Ipv6 unicast-routing
!
Router ospfv3 4
    Router-id 192.168.1.1
  !
Interface E 0/0
 Ipv6 enable
 Ip address 10.1.1.1 255.255.255.0
 Ospfv3 4 area 0 ipv4
 No shut
!
Interface Loopback0
 Ipv6 enable
 Ipv4 172.16.1.1 255.255.255.0
 Ospfv3 4 area 0 ipv4
```

Refer to the exhibit. The network administrator configured the branch router for IPv6 on the E 0/0 interface. The neighboring router is fully configured to meet requirements, but the neighbor relationship is not coming up.

Which action fixes the problem on the branch router to bring the IPv6 neighbors up?

A. Disable OSPF for IPv4 using the no ospfv3 4 area 0 ipv4 command under the E 0/0 interface.

B. Enable the IPv4 address family under the router ospfv3 4 process by using the address-family ipv4 unicast command.

C. Disable IPv6 on the E 0/0 interface using the no ipv6 enable command.

D. Enable the IPv4 address family under the E 0/0 interface by using the address-family ipv4 unicast command.

Correct Answer: *B*

Community vote distribution

B (100%)

Ll123123 1 month, 3 weeks ago

Selected Answer: B

i choose B
because OSPFv3 configuration require the ipv4 and ipv6 address family configuration under the ospfv3 process.

upvoted 1 times

Reikidude00 1 year, 5 months ago

Selected Answer: B

B is correct here, af configuration must be done under router-config

```
R1#show running-config | section router
router ospfv3 4
router-id 1.1.1.1
!
address-family ipv4 unicast
exit-address-family
```

upvoted 1 times

Dacusai 1 year, 5 months ago

I think the question is wrong from the beginning, The appropriate address family is enabled automatically, but at least you should have a link local address for the relation to form.

upvoted 2 times

Networkingguy 1 year, 10 months ago

Selected Answer: B

B is the correct answer here

upvoted 1 times

wts 1 year, 11 months ago

Doesn't "address-family ipv4 unicast" automatically appear under the router after adding an interface to OSPF?

upvoted 2 times

wts 1 year, 10 months ago

"The appropriate address family is enabled automatically when OSPFv3 is enabled on an interface.", - cisco official guide.

upvoted 2 times

wts 1 year, 10 months ago

```
(config)#ipv6 unicast-routing
(config)#ipv6 cef

(config)#router ospfv3 4
(config-router)#router-id 192.168.1.1

(config)#interface GigabitEthernet0/0
(config-if)#ipv6 enable
(config-if)#ipv6 ospf neighbor FE80: - this and the previous commands are needed because the ipv6 address is not specified.
(config-if)#ospfv3 4 ipv6 area 0 - this command is missing so that the interface participates in ospfv3 ipv6. (this is the correct order of the keywords, the wrong order is given in the question.)
```

To be honest, I do not know what to do if I get a question like this.

upvoted 3 times

Pietjeplukgeluk 1 month, 1 week ago

There is no need to statically configure a neighbor for IPv6 on a broadcast network type.
Just tested, and it brings up OSPF neighbor (note there is no need to configure any routable IPv6 address for just OSPF neighbor on the interface)

```
ipv6 unicast-routing
!
interface GigabitEthernet0/0
ipv6 enable
5
ospfv3 1 ipv6 area 0
!
!
router ospfv3 1
router-id 1.1.1.1
!
address-family ipv6 unicast
exit-address-family
!
!
router ospfv3 1
router-id 1.1.1.1
!
address-family ipv6 unicast
exit-address-family
!
!
router ospfv3 1
router-id 1.1.1.1
!
address-family ipv6 unicast
exit-address-family
!
```

Note the " address-family ipv6 unicast" is added by default.

upvoted 1 times

Azaelyus 11 months, 2 weeks ago

call Miroslav Tihlarik

upvoted 1 times

Azaelyus 11 months, 2 weeks ago

+420725950480

upvoted 1 times

myrmike 2 years ago

spapi0390 is right on. Looking at the config of the interfaces only ospfv3 ipv4 is enabled on the interfaces. The presumption being that the neighbor router is configured for ospfve ipv4 and not ipv6

upvoted 1 times

Jenia1 2 years, 1 month ago

B is the correct answer, but I totally agree with amgue, the question is about IPv6 and not IPv4, although if you configure OSPFv3 like this, you will see LSA type 8/9, so maybe what they meant, but it is interesting why they call it IPv6 as you also will see LSA type 1/2.
Appreciate if someone can shed a light on this

upvoted 3 times

spapi0390 2 years, 1 month ago

With the OSPFv3 address families feature, you may have two device processes per interface, but only one process per AF. If the IPv4 AF is used, an IPv4 address must first be configured on the interface. For IPv6 AF it is enough, if only IPv6 is enabled on the interface, as OSPFv3 uses link-local addresses. A single IPv4 or IPv6 OSPFv3 process running multiple instances on the same interface is not supported.

upvoted 1 times

amgue 2 years, 1 month ago

I don't understand the answer ! the question is about the enabling Ipv6 neighborship, not IPV4 neighborship, can somewone explain this to me PLEASE ?

upvoted 3 times

spapi0390 2 years, 1 month ago

With the OSPFv3 address families feature, you may have two device processes per interface, but only one process per AF. If the IPv4 AF is used, an IPv4 address must first be configured on the interface. For IPv6 AF it is enough, if only IPv6 is enabled on the interface, as OSPFv3 uses link-local addresses. A single IPv4 or IPv6 OSPFv3 process running multiple instances on the same interface is not supported.

upvoted 4 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct
router ospfv3 [process-id]

address-family ipv4 unicast

upvoted 1 times

Question #29

Topic 1

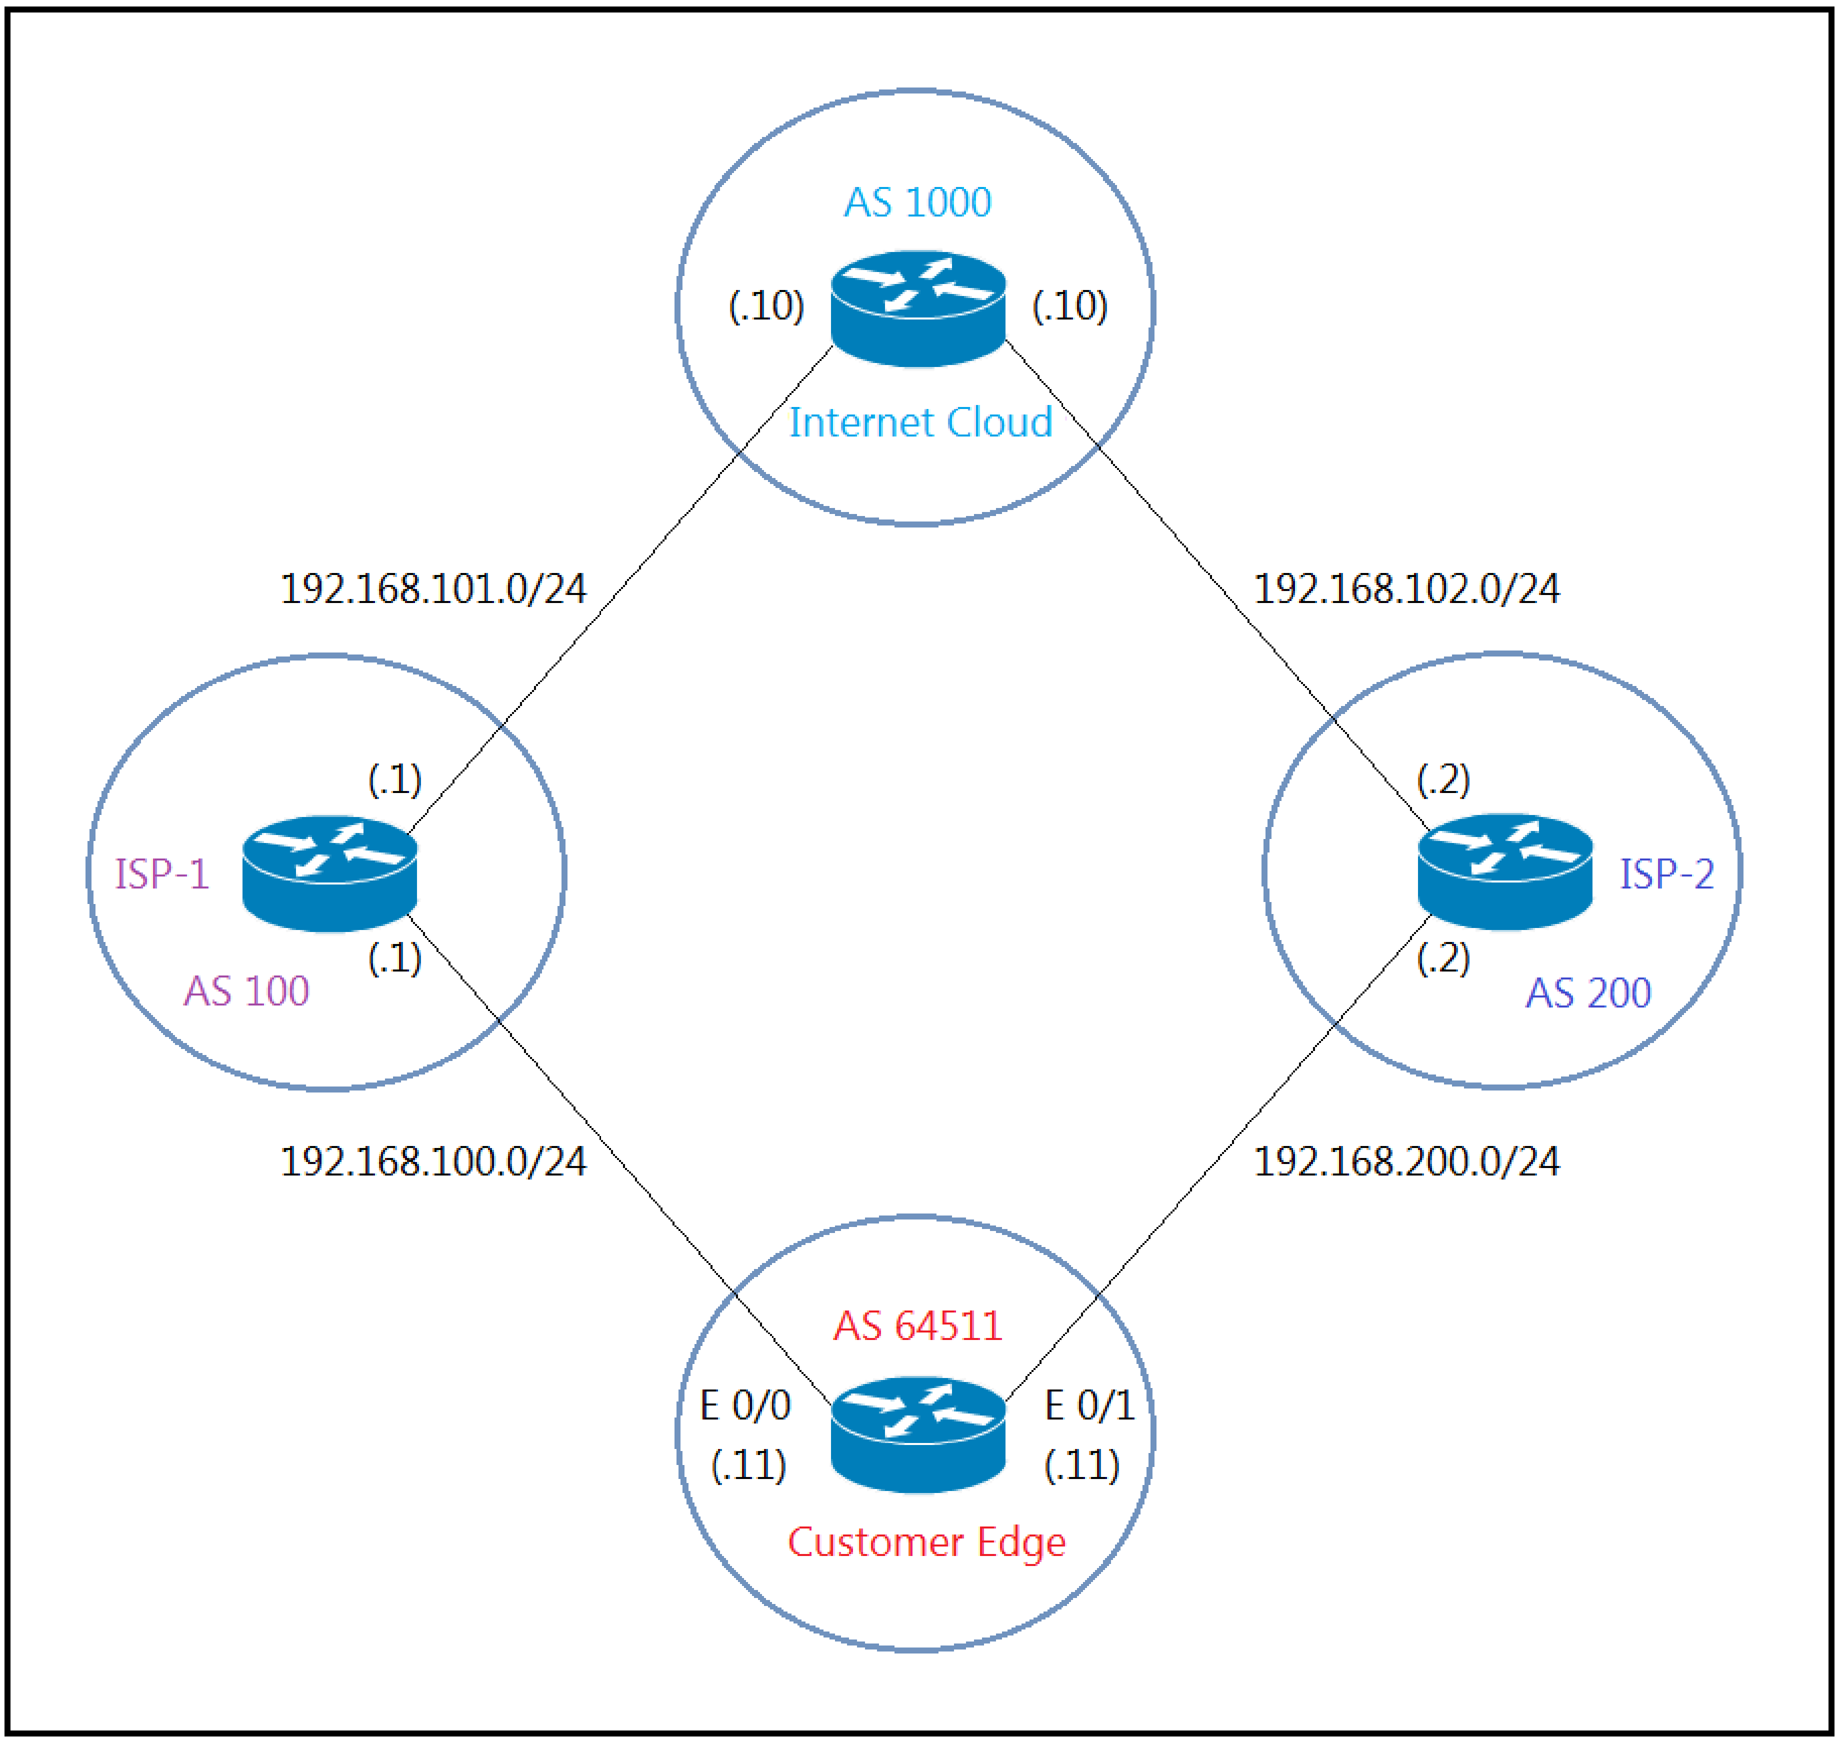

Refer to the exhibit. The network administrator has configured the Customer Edge router (AS 64511) to send only summarized routes toward ISP-1 (AS 100) and
ISP-2 (AS 200).

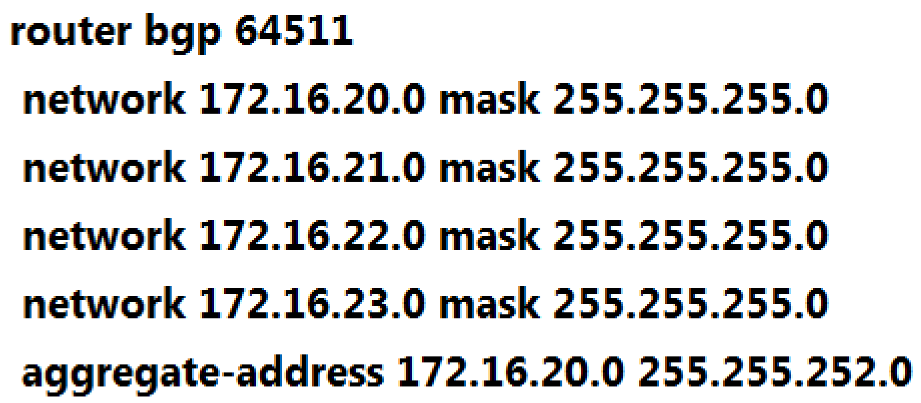

```
router bgp 64511
 network 172.16.20.0 mask 255.255.255.0
 network 172.16.21.0 mask 255.255.255.0
 network 172.16.22.0 mask 255.255.255.0
 network 172.16.23.0 mask 255.255.255.0
 aggregate-address 172.16.20.0 255.255.252.0
```

After this configuration, ISP-1 and ISP-2 continue to receive the specific routes and the summary route.

Which configuration resolves the issue?

A.

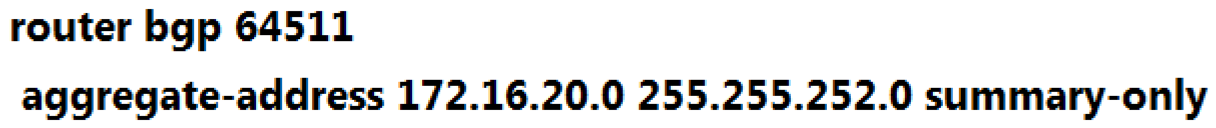

```
router bgp 64511
 aggregate-address 172.16.20.0 255.255.252.0 summary-only
```

B.

```
router bgp 64511
 neighbor 192.168.100.1 summary-only
 neighbor 192.168.200.2 summary-only
```

C.

```
ip prefix-list PL_BLOCK_SPECIFIC deny 172.16.20.0/22 ge 22
ip prefix-list PL_BLOCK_SPECIFIC permit 172.16.20.0/22
!
route-map BLOCK_SPECIFIC permit 10
 match ip address prefix-list PL_BLOCK_SPECIFIC
!
router bgp 64511
 aggregate-address 172.16.20.0 255.255.252.0 suppress-map BLOCK_SPECIFIC
```

D.

```
interface E 0/0
 ip bgp suppress-map BLOCK_SPECIFIC
!
interface E 0/1
 ip bgp suppress-map BLOCK_SPECIFIC
!
 ip prefix-list PL_BLOCK_SPECIFIC permit 172.16.20.0/22 ge 24
!
route-map BLOCK_SPECIFIC permit 10
 match ip address prefix-list PL_BLOCK_SPECIFIC
```

Correct Answer: *A*

Huntkey 1 year, 3 months ago

A is correct

upvoted 2 times

Dataset 1 year, 4 months ago

Correct A

upvoted 2 times

timtgh 1 year, 6 months ago

The aggregate-address command causes a router to send a summary route AND still also send the normal individual routes - unless you add "summary-only." So A is correct.

upvoted 1 times

AliMo123 2 years, 4 months ago

A is the correct answer but let me clarify this:
aggregate address is similar to route summary addresses like a router prefers to receive a single route instead of too many.
notice here 172.16.20.0 255.255.252.0 covers from 172.16.21.0 to 172.16.23.0

so aggregate-add 172.16.20.0 255.255.252.0 summary only will force these addresses to send only summary to AS 100 and 200 to fulfill the request

upvoted 3 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

youarewrong 2 years, 5 months ago

Set this up to verify. Answer 'A' does indeed summarize on ISP-1 and 2.

upvoted 1 times

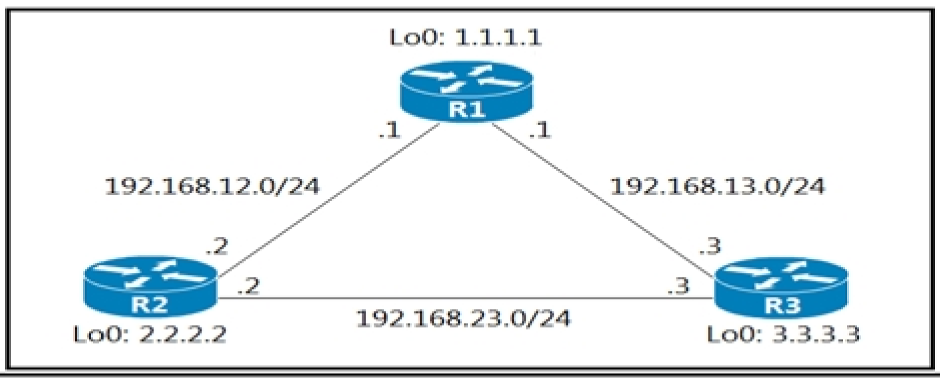

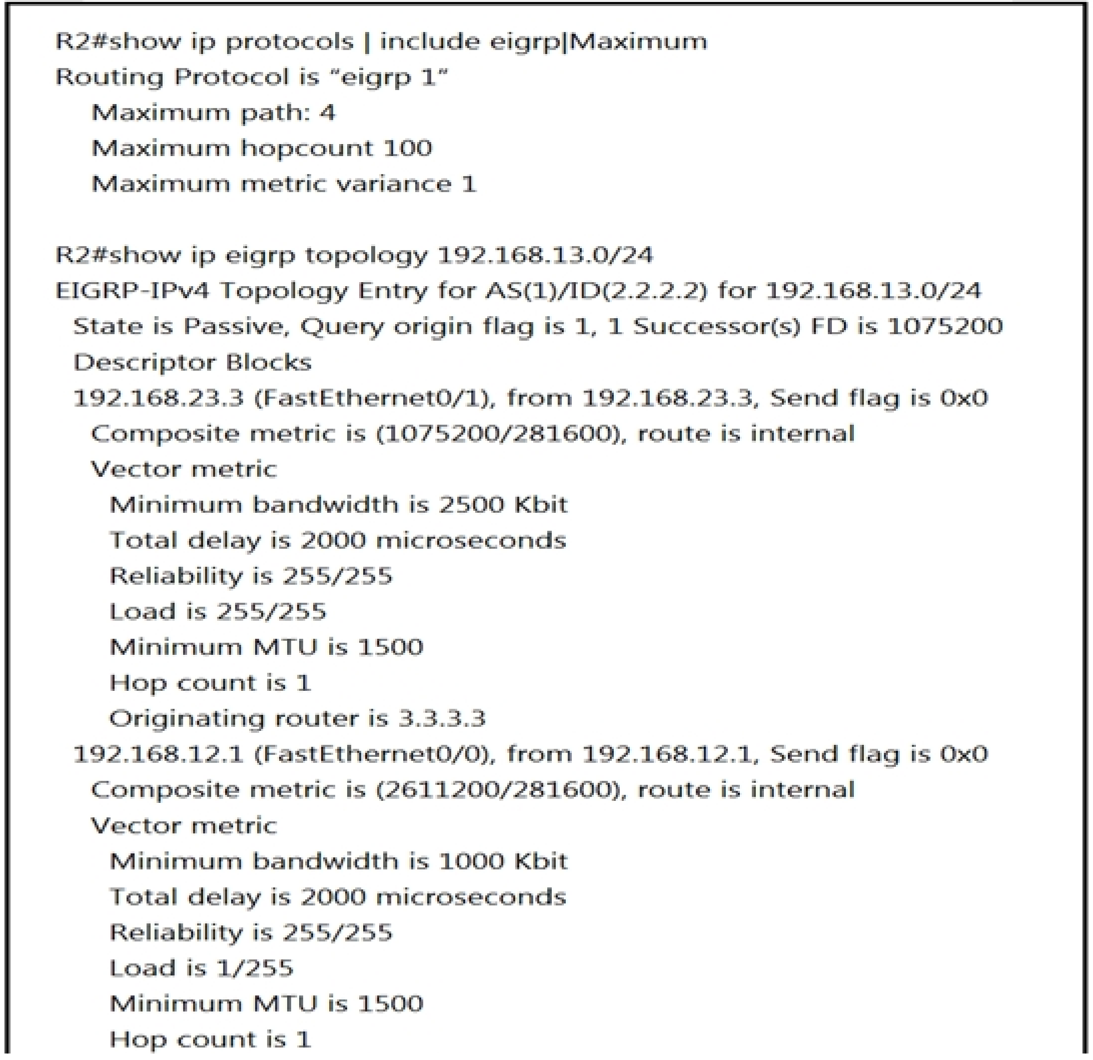

```
R2#show ip protocols | include eigrp|Maximum
Routing Protocol is "eigrp 1"
    Maximum path: 4
    Maximum hopcount 100
    Maximum metric variance 1

R2#show ip eigrp topology 192.168.13.0/24
EIGRP-IPv4 Topology Entry for AS(1)/ID(2.2.2.2) for 192.168.13.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s) FD is 1075200
  Descriptor Blocks
  192.168.23.3 (FastEthernet0/1), from 192.168.23.3, Send flag is 0x0
    Composite metric is (1075200/281600), route is internal
    Vector metric
      Minimum bandwidth is 2500 Kbit
      Total delay is 2000 microseconds
      Reliability is 255/255
      Load is 255/255
      Minimum MTU is 1500
      Hop count is 1
      Originating router is 3.3.3.3
  192.168.12.1 (FastEthernet0/0), from 192.168.12.1, Send flag is 0x0
    Composite metric is (2611200/281600), route is internal
    Vector metric
      Minimum bandwidth is 1000 Kbit
      Total delay is 2000 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 1
```

Refer to the exhibit. R2 has two paths to reach 192.168.13.0/24, but traffic is sent only through R3.

Which action allows traffic to use both paths?

A. Configure the variance 4 command under the EIGRP process on R2.

B. Configure the bandwidth 2000 command under interface FastEthernet0/0 on R2.

C. Configure the delay 1 command under interface FastEthernet0/0 on R2.

D. Configure the variance 2 command under the EIGRP process on R2.

Correct Answer: *A*

Community vote distribution

A (100%)

TECH3K3 Highly Voted 1 year, 10 months ago

Selected Answer: A

Answer = A
Explanation:
Feasible Distance of R3 (successor) = 1075200
Feasible Distance of R1 (feasible successor) = 2611200

Calculation: 2611200 / 1075200 = 2.4

So we need a Variance value higher than 2.4 for unequal cost load balancing to work.

upvoted 15 times

xziomal9 Most Recent 1 year, 8 months ago

Selected Answer: A

The correct answer is: A

upvoted 1 times

davdtech 1 year, 10 months ago

I do not understand
The successor has an FD of 1075200
The feasible successor has an RD of 281600 so the feasibility condition is met ...why should we manipulate the variance? maybe there is a missing zero in the RD of the Feasible successor?

upvoted 1 times

timtgh 1 year, 6 months ago

Feasibility isn't the issue. R3 is the BEST route, so it's the only route used. Unless we change the variance. With variance 4, we can use any route where the FD is not greater than 4 times 1075200. (Variance 2 wouldn't be high enough in this case.)

upvoted 4 times

weltongama 1 year, 10 months ago

Selected Answer: A

The given answer is correct

upvoted 1 times

Surfside92 2 years, 1 month ago

The reason the given answer A is correct and why its not B.
If we want to enable load balancing we have to use the following formula:
FD of feasible successor < FD of successor * multiplier
So we can work out that FD of feasible successor (R2) / FD of successor (R3) = 2.4
So the multiplier - or variance needs to be more than 2.4 - which means only answer A is correct.

upvoted 2 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #31 Topic 1

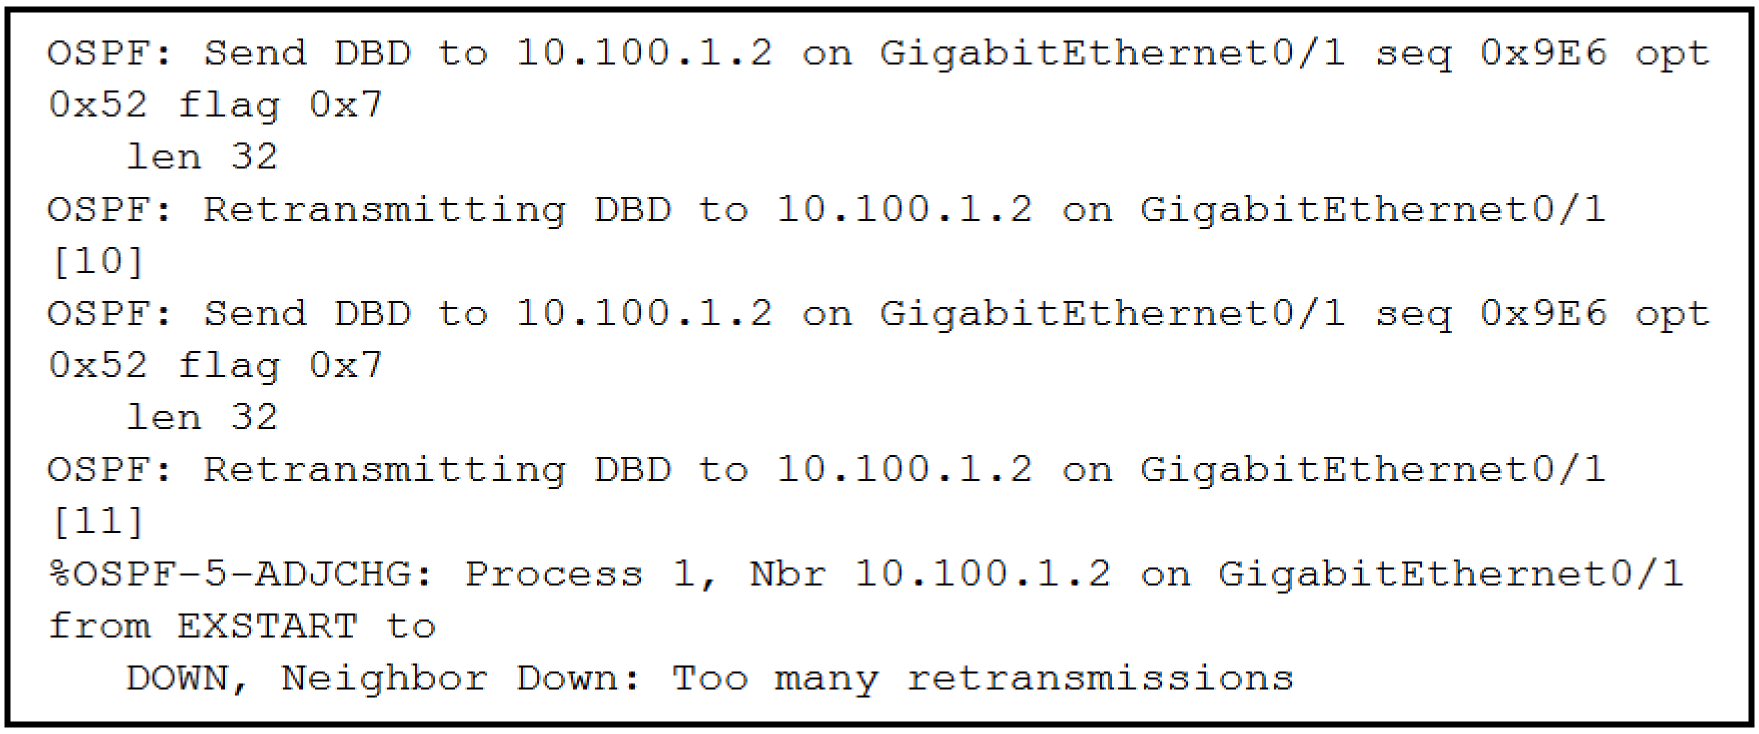

```
OSPF: Send DBD to 10.100.1.2 on GigabitEthernet0/1 seq 0x9E6 opt
0x52 flag 0x7
   len 32
OSPF: Retransmitting DBD to 10.100.1.2 on GigabitEthernet0/1
[10]
OSPF: Send DBD to 10.100.1.2 on GigabitEthernet0/1 seq 0x9E6 opt
0x52 flag 0x7
   len 32
OSPF: Retransmitting DBD to 10.100.1.2 on GigabitEthernet0/1
[11]
%OSPF-5-ADJCHG: Process 1, Nbr 10.100.1.2 on GigabitEthernet0/1
from EXSTART to
   DOWN, Neighbor Down: Too many retransmissions
```

Refer to the exhibit. The OSPF neighbor relationship is not coming up.

What must be configured to restore OSPF neighbor adjacency?

A. matching hello timers

B. OSPF on the remote router

C. use router ID

D. matching mtu values

Correct Answer: *D*

Community vote distribution

D (100%)

Hack4 1 year, 10 months ago

The given answer is correct

upvoted 1 times

Networkingguy 1 year, 10 months ago

Selected Answer: D

D is correct here

upvoted 1 times

Girmiti 1 year, 11 months ago

Selected Answer: D

```
*Jan 10 07:12:55.724: OSPF-1 ADJ Et0/0: Rcv DBD from 10.10.10.2 seq 0x1BFC opt 0x52 flag 0x7 len 32 mtu 100 state EXCHANGE
*Jan 10 07:12:55.724: OSPF-1 ADJ Et0/0: Nbr 10.10.10.2 has smaller interface MTU
*Jan 10 07:12:55.724: OSPF-1 ADJ Et0/0: Send DBD to 10.10.10.2 seq 0x1BFC opt 0x52 flag 0x0 len 32
IOU1#undebug all
All possible debugging has been turned off
IOU1#
IOU1#
IOU1#
*Jan 10 07:13:06.410: %OSPF-5-ADJCHG: Process 1, Nbr 10.10.10.2 on Ethernet0/0 from EXCHANGE to DOWN, Neighbor Down: Too many retransmissions
IOU1#
*Jan 10 07:14:06.417: %OSPF-5-ADJCHG: Process 1, Nbr 10.10.10.2 on Ethernet0/0 from DOWN to DOWN, Neighbor Down: Ignore timer expired
```

upvoted 2 times

Surfside92 2 years, 1 month ago

D is correct :
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13684-12.html

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

AliMo123 2 years, 4 months ago

D is correct
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/119384-technote-ospf-00.html

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #32 Topic 1

An engineer configured two routers connected to two different service providers using BGP with default attributes. One of the links is presenting high delay, which causes slowness in the network.
Which BGP attribute must the engineer configure to avoid using the high-delay ISP link if the second ISP link is up?

A. AS-PATH

B. WEIGHT

C. MED

D. LOCAL_PREF

Correct Answer: *D*

Community vote distribution

D (100%)

JingleJangus Highly Voted 1 year, 11 months ago

Selected Answer: D

It D. The key word is TWO routers in the AS. We need a solution that works for both, and only needs to be configured once. Yes, weight could work, however it isnt the best solution here since its locally significant.

upvoted 5 times

Ll123123 Most Recent 1 month, 3 weeks ago

Selected Answer: D

I choose D. Because the question said he is configuring the "2" router connecting to ISP, if he can configure the router connecting to the two routers, he can change the WEIGHT. But if he is on the two router, the only config that can influence other is Local_Pref.

upvoted 1 times

SnoopDD 2 months ago

i think it's A

upvoted 1 times

CosmasNyoni 5 months, 3 weeks ago

I choose D

upvoted 1 times

Huntkey 1 year, 3 months ago

I think it is A. only A is possible to affect choosing for both directions. The other options only affect one direction.

upvoted 2 times

Huntkey 1 year, 3 months ago

Can you say you avoided using one link if only one direction is not using the link but the other direction continues to use it?

upvoted 1 times

Spyrous 1 year, 9 months ago

If the question stated TWO cisco routers, then the answer would be B. Setting the weight with an inbound route map, you can influence outbound traffic. The same goes for local preference, second attribute in the list of preference in cisco routers, but first in non-cisco routers. I agree that the answer should be D.

upvoted 1 times

Macferson 1 year, 10 months ago

The answer is D,
The Local Preference attribute is used:
To choose between multiple exit paths from your network

- Non-transitive and optional attribute
- Local to an AS only
* Default local preference is 100 (IOS)
- Used to influence BGP path selection
* Determines the best path for outbound traffic
- Path with highest local preference wins

upvoted 4 times

Raider1 2 years, 2 months ago

D is correct:
Three steps are by far the most important ones.

Prefer the path with the highest local preference

Prefer the path with the shortest AS path
Prefer the path with the lowest multi-exit discriminator (MED)
upvoted 2 times

error_909 2 years, 3 months ago
The given answer is correct
upvoted 1 times

geek1992 2 years, 4 months ago
D IS CORRECT
upvoted 1 times

geek1992 2 years, 4 months ago
A and 100% sure
upvoted 2 times

cryptonite 2 years, 3 months ago
AS_PATH prepends allows you to ensure that returning traffic to the network uses path with the short AS-Path. I will go with AS-Path because with it, one can influence how traffic returns from the rest of the world.

Local preference only determines how traffic exits the network if there are multiple routers so this is wrong.

MED is like metric, MED is what we learn from our upstream service provider or we set the Med on the way in. But Med only chooses the exit link not how the traffic returns.
upvoted 1 times

cryptonite 2 years, 3 months ago
Local Preference is correct, apologies. There are multiple routers.
upvoted 1 times

studybuddy10 2 years, 1 month ago
The production world answer for this is both A and D, prepend adverts out the slower link and use local preference to exit out the better link. The cisco answer is local preference as they mention 2 routes which is the key here.
upvoted 2 times

Hammad745 2 years, 4 months ago
Weight is the correct answer
upvoted 1 times

vdsdrs 2 years, 4 months ago
No, weight is local to router. In the question we have two routers connected to two ISPs.
Local_pref is exchanged between IBGP peers.
D is correct answer.
upvoted 1 times

AliMo123 2 years, 4 months ago
D is correct
You can use local preference to choose the outbound external BGP path.
upvoted 1 times

examShark 2 years, 4 months ago
B weight
upvoted 1 times

examShark 2 years, 4 months ago
The given answer is correct
upvoted 1 times

cyrus777 1 year, 8 months ago
weight is local to router. In the question we have two routers connected to two ISPs.
Local_pref is exchanged between IBGP peers.
D is correct answer.
upvoted 2 times

Question #33 Topic 1

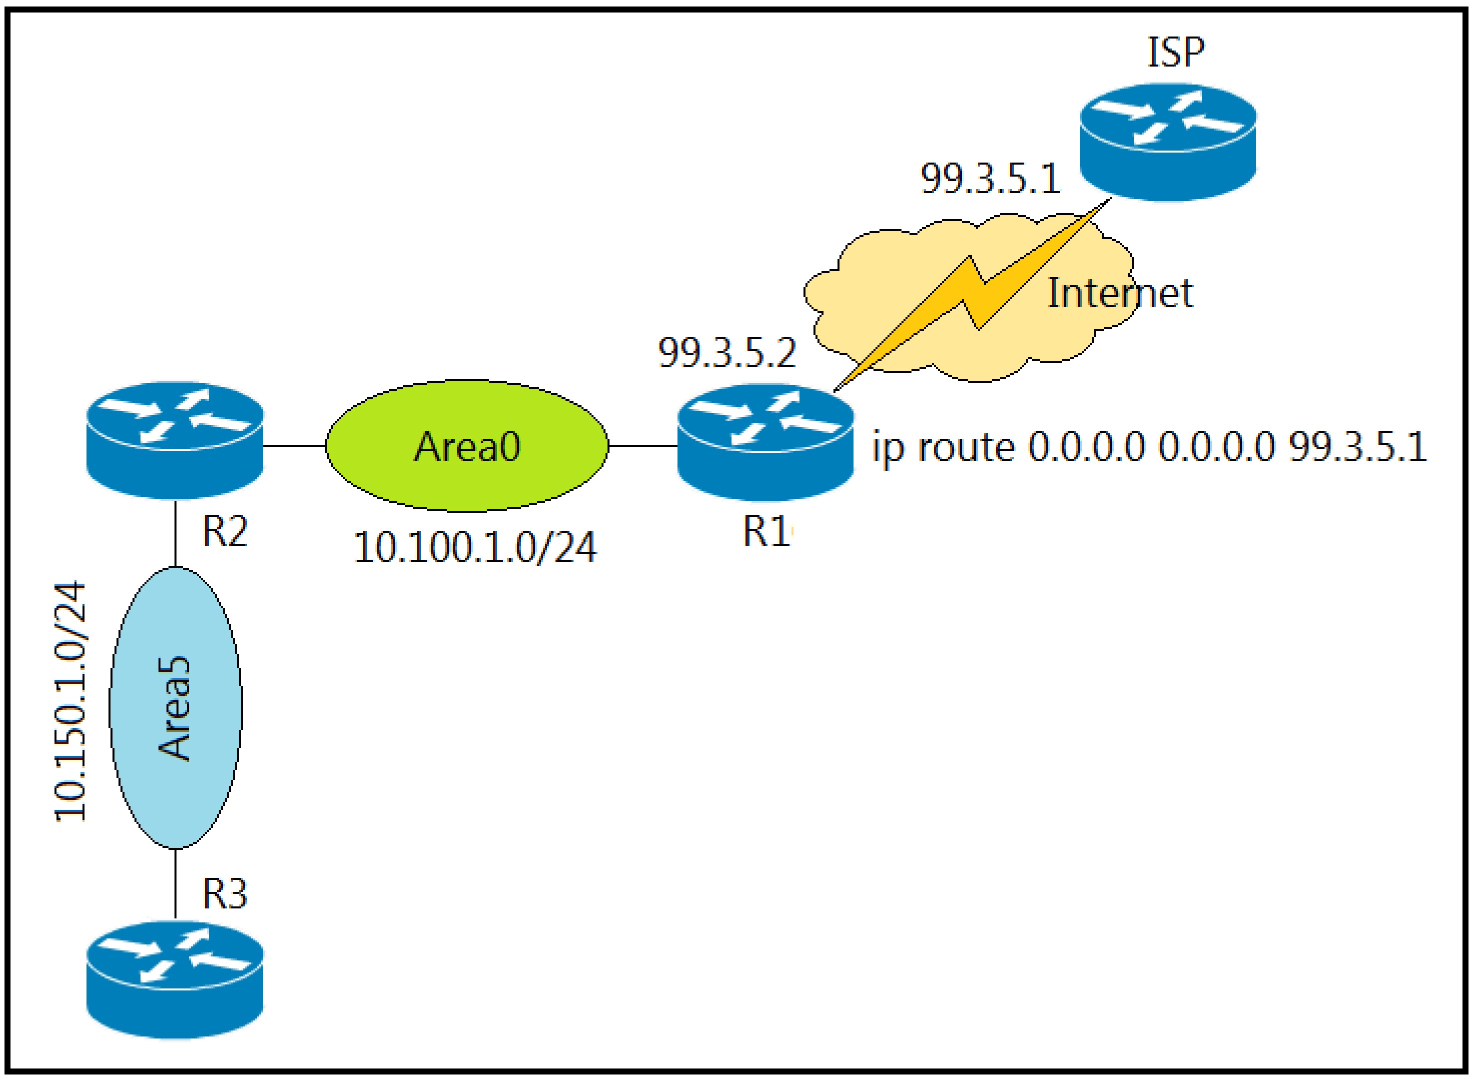

Refer to the exhibit. A network administrator redistributed the default static route into OSPF toward all internal routers to reach to Internet. Which set of commands restores reachability to the Internet by internal routers?

A. router ospf 1 redistribute static subnets

B. router ospf 1 network 0.0.0.0 0.0.0.0 area 0

C. router ospf 1 redistribute connected 0.0.0.0

D. router ospf 1 default-information originate

Correct Answer: *D*

⊟ **Dimma** Highly Voted 1 year, 11 months ago

Why do we need to do in ospf default-information originate always ? As you probably already know, default-information originate tells the router to inject any default route that has been configured on the router into the OSPF. The OSPF router does not, by default, generate a default route into the OSPF domain.
In OSPF, the "default-information originate" command will not advertise to any other routers without a default route in the routing table. When added the "always" keyword , it tells the router to advertise a default route to other routers even if you don't have a default route in the routing table

upvoted 14 times

⊟ **Networkingguy** 1 year, 11 months ago

Spot on Dimma, upvoted. Now if you can stop cheating on your wife and bring Richmond back into Premiershipform, that would be hugely appreciated.

upvoted 4 times

⊟ **Hack4** Most Recent 1 year, 10 months ago

the given answer is correct

upvoted 1 times

⊟ **error_909** 2 years, 3 months ago

The given answer is correct

upvoted 1 times

⊟ **examShark** 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #34 *Topic 1*

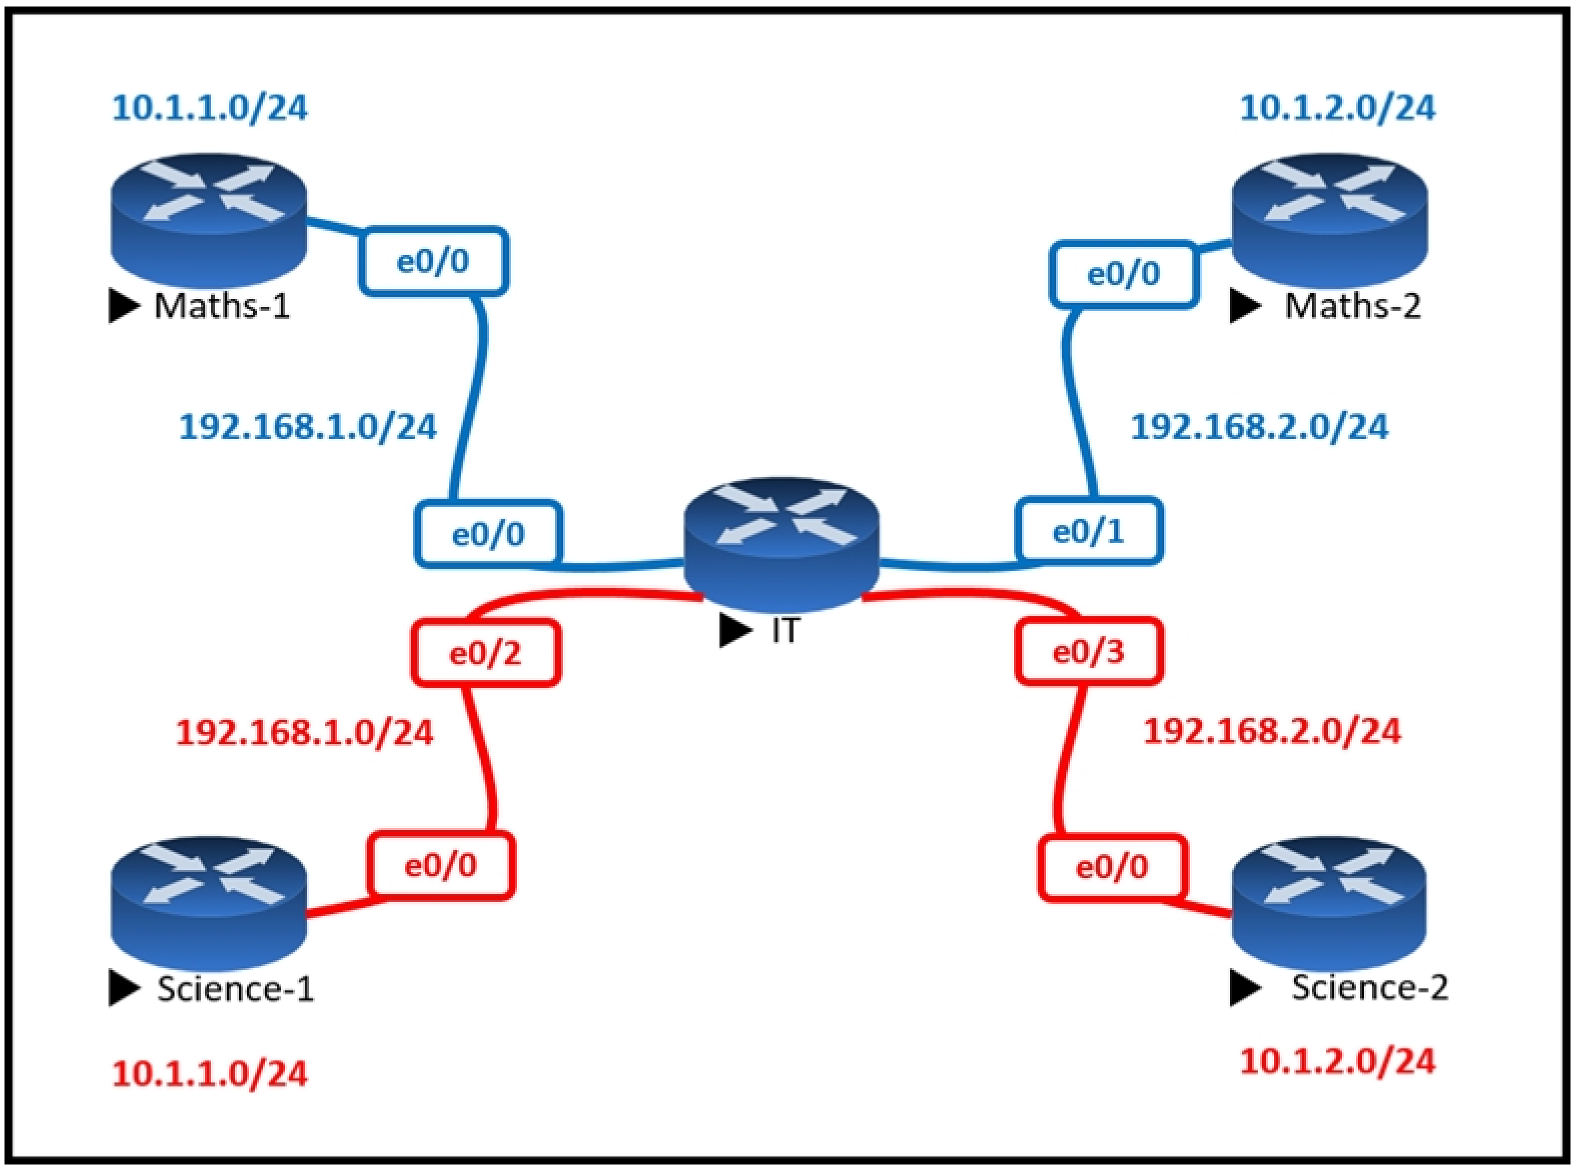

Refer to the exhibit. The Math and Science departments connect through the corporate IT router, but users in the Math department must not be able to reach the
Science department and vice versa.
Which configuration accomplishes this task?

A. vrf definition Science address-family ipv4 ! interface E 0/2 ip address 192.168.1.1 255.255.255.0 no shut ! interface E 0/3 ip address 192.168.2.1 255.255.255.0 no shut

B. vrf definition Science address-family ipv4 ! interface E 0/2 vrf forwarding Science ip address 192.168.1.1 255.255.255.0 no shut ! interface E 0/3 vrf forwarding Science ip address 192.168.2.1 255.255.255.0 no shut

C. vrf definition Science address-family ipv4 ! interface E 0/2 ip address 192.168.1.1 255.255.255.0 vrf forwarding Science no shut ! interface E 0/3 ip address 192.168.2.1 255.255.255.0 vrf forwarding Science no shut

D. vrf definition Science ! interface E 0/2 ip address 192.168.1.1 255.255.255.0 no shut ! interface E 0/3 ip address 192.168.2.1 255.255.255.0 no shut

Correct Answer: *B*

Community vote distribution

B (100%)

bjromero28 Highly Voted 2 years, 2 months ago

Answer B is correct.

Remember that you must add the vrf to the interface first and then the ip address. Adding the ip address before the vrf forwarding will remove the ip address from the interface.

upvoted 9 times

Carl1999 Highly Voted 1 year, 10 months ago

A.
vrf definition Science
address-family ipv4
!
interface E 0/2
ip address 192.168.1.1 255.255.255.0
no shut!
interface E 0/3
ip address 192.168.2.1 255.255.255.0

```
no shut
```

B.

```
vrf definition Science
address-family ipv4
!
interface E 0/2
vrf forwarding Science
ip address 192.168.1.1 255.255.255.0
no shut
!
interface E 0/3
vrf forwarding Science
ip address 192.168.2.1 255.255.255.0
no shut
```

C.

```
vrf definition Science
address-family ipv4
!
interface E 0/2
ip address 192.168.1.1 255.255.255.0
vrf forwarding Science
no shut
!
interface E 0/3
ip address 192.168.2.1 255.255.255.0
vrf forwarding Science
no shut
```

D.

```
vrf definition Science
!
interface E 0/2
ip address 192.168.1.1 255.255.255.0
no shut
!
interface E 0/3
ip address 192.168.2.1 255.255.255.0
no shut
```

upvoted 7 times

Noproblem22 Most Recent 1 year, 1 month ago

B is correct

upvoted 1 times

Hack4 1 year, 10 months ago

Yes i agree. B is correct

upvoted 1 times

wts 1 year, 10 months ago

Why are the answer options in a line? ... you can also write in light gray small print.

upvoted 1 times

wts 1 year, 10 months ago

```
ip vrf Science
interface e0/2
ip vrf forwarding Science
ip address 192.168.1.1 255.255.255.0
interface e0/2
ip vrf forwarding Science
ip address 192.168.2.1 255.255.255.0
```

Entering the vrf-bound address-family ipv4 configuration mode is typically used to configure BGP. It's not clear why it's here.

upvoted 1 times

Girmiti 1 year, 11 months ago

Selected Answer: B

B is correct.

upvoted 1 times

Question #35 *Topic 1*

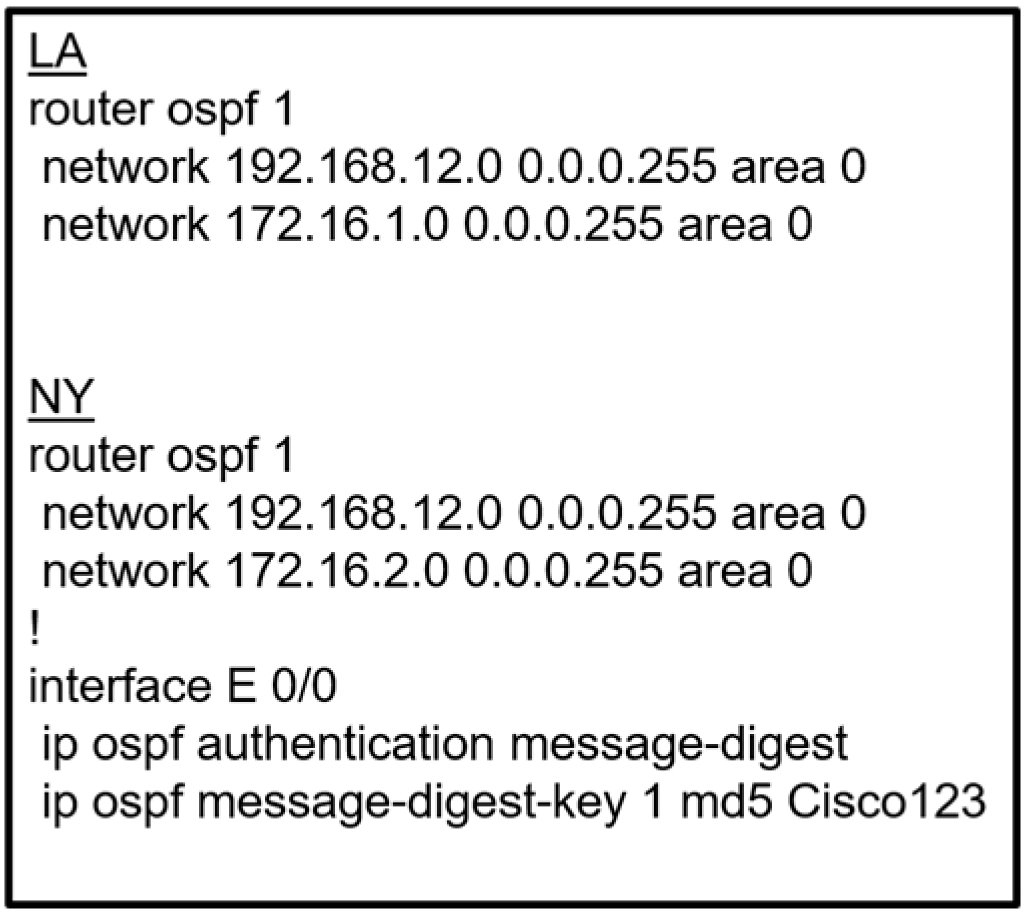

```
LA
router ospf 1
 network 192.168.12.0 0.0.0.255 area 0
 network 172.16.1.0 0.0.0.255 area 0


NY
router ospf 1
 network 192.168.12.0 0.0.0.255 area 0
 network 172.16.2.0 0.0.0.255 area 0
!
interface E 0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 Cisco123
```

Refer to the exhibit. The neighbor relationship is not coming up.
Which two configurations bring the adjacency up? (Choose two.)

A. LA interface E 0/0 ip ospf authentication-key Cisco123

B. NY interface E 0/0 no ip ospf message-digest-key 1 md5 Cisco123 ip ospf authentication-key Cisco123

C. LA interface E 0/0 ip ospf message-digest-key 1 md5 Cisco123

D. LA router ospf 1 area 0 authentication message-digest

E. NY router ospf 1 area 0 authentication message-digest

Correct Answer: *CD*

Community vote distribution

CD (100%)

Malasxd 7 months, 2 weeks ago

Selected Answer: CD

C and D are correct

upvoted 1 times

studybuddy10 2 years, 1 month ago

Agree with C and D, tested exact configurations.

upvoted 1 times

AliMo123 2 years, 1 month ago

C and D are correct
remember: one of the rule for OSPF neighbor relationship to come up is matching Authentications

upvoted 1 times

Question #36 Topic 1

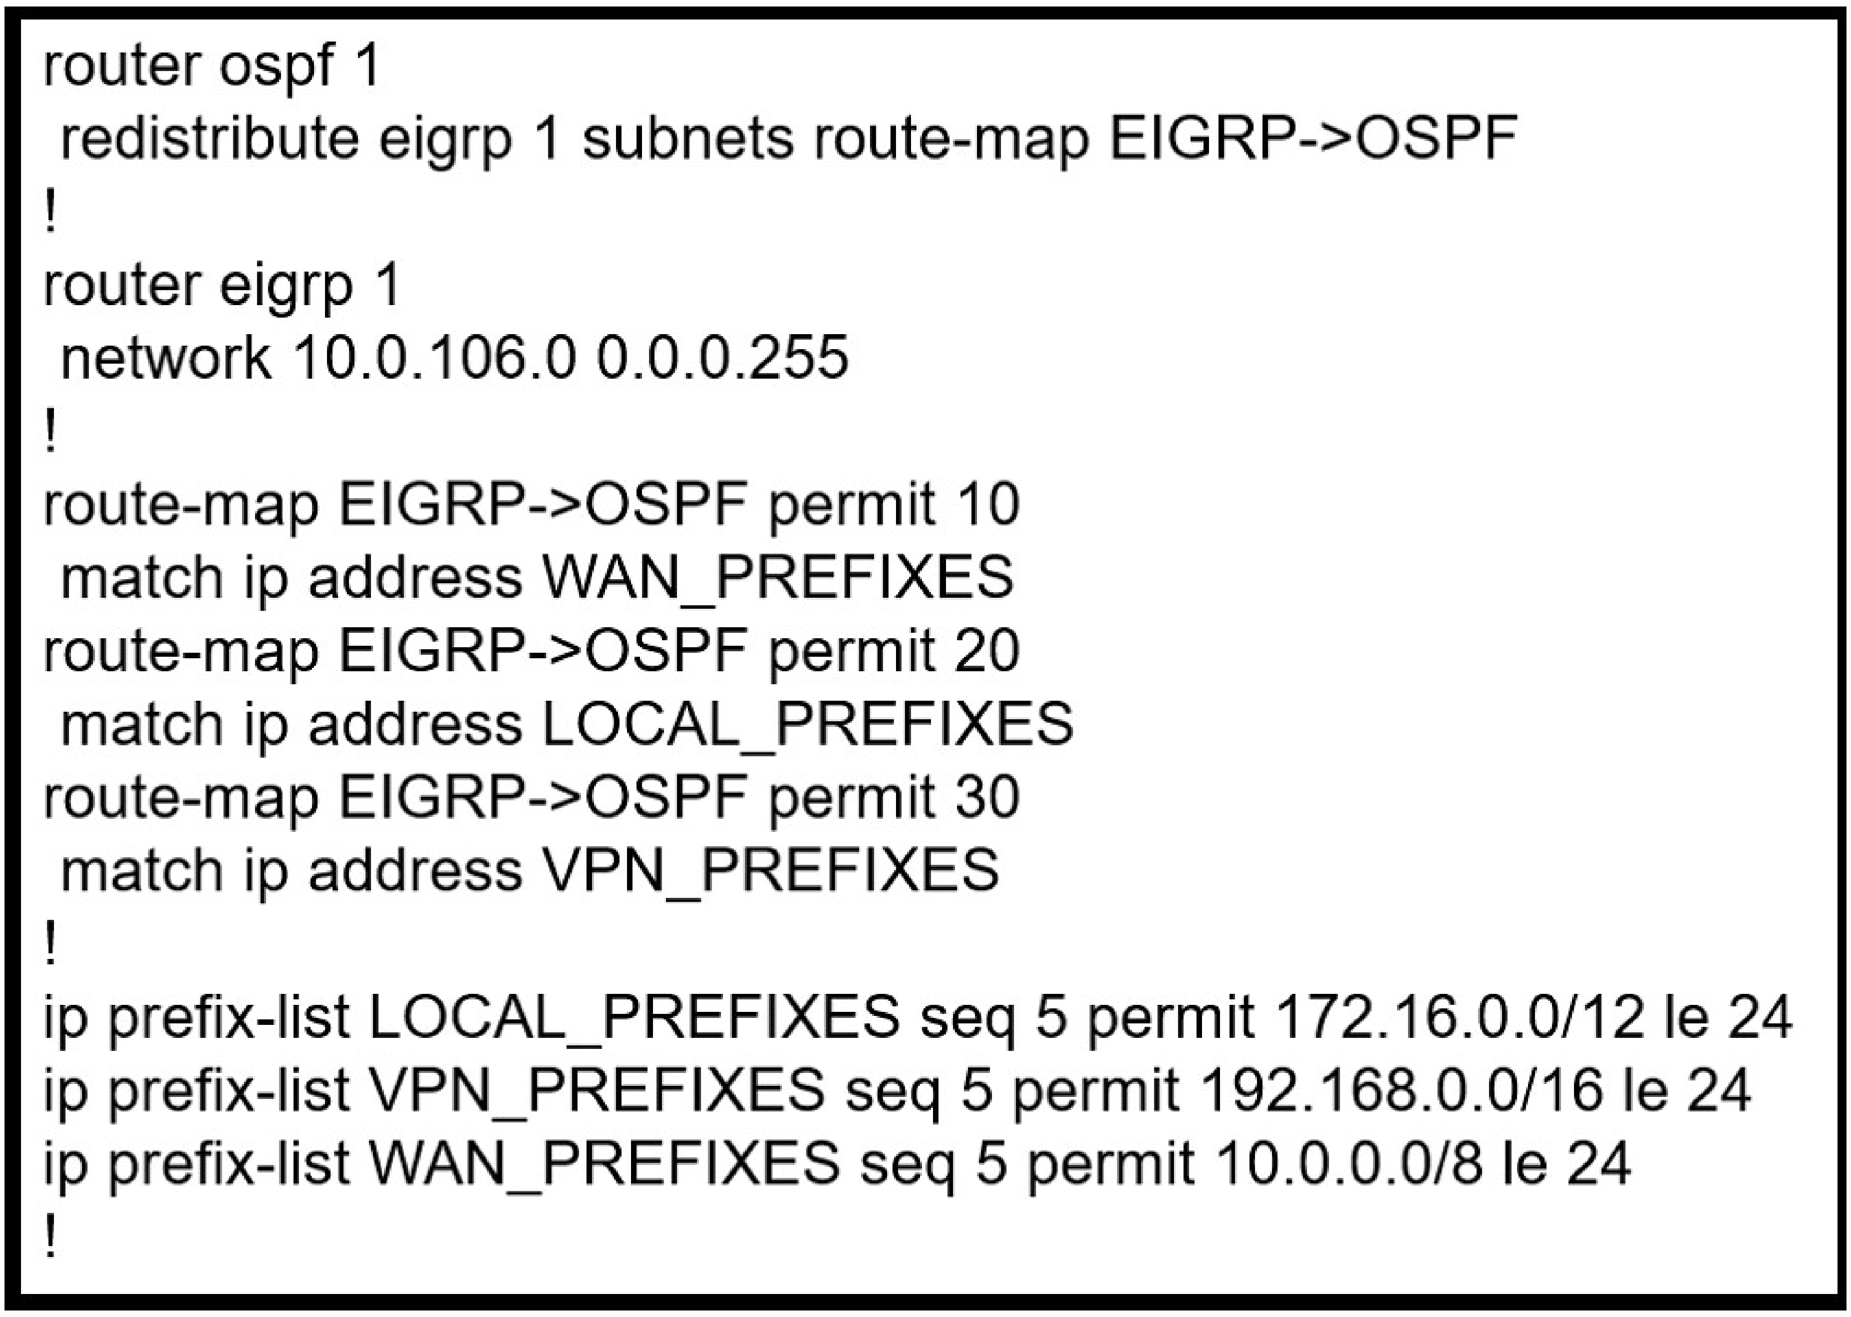

```
router ospf 1
 redistribute eigrp 1 subnets route-map EIGRP->OSPF
!
router eigrp 1
 network 10.0.106.0 0.0.0.255
!
route-map EIGRP->OSPF permit 10
 match ip address WAN_PREFIXES
route-map EIGRP->OSPF permit 20
 match ip address LOCAL_PREFIXES
route-map EIGRP->OSPF permit 30
 match ip address VPN_PREFIXES
!
ip prefix-list LOCAL_PREFIXES seq 5 permit 172.16.0.0/12 le 24
ip prefix-list VPN_PREFIXES seq 5 permit 192.168.0.0/16 le 24
ip prefix-list WAN_PREFIXES seq 5 permit 10.0.0.0/8 le 24
!
```

Refer to the exhibit. The network administrator configured redistribution on an ASBR to reach to all WAN networks but failed.
Which action resolves the issue?

A. The route map EIGRP->OSPF must have the 10.0.106.0/24 entry to exist in one of the three prefix lists to pass

B. EIGRP must redistribute the 10.0.106.0/24 route instead of using the network statement

C. The OSPF process must have a metric when redistributing prefixes from EIGRP

D. The route map must have the keyword prefix-list to evaluate the prefix list entries

Correct Answer: *D*

Community vote distribution

D (100%)

Brand 3 months, 4 weeks ago

Selected Answer: D

```
R1(config-route-map)#match ip address ?
<1-199> IP access-list number
<1300-2699> IP access-list number (expanded range)
WORD IP access-list name
prefix-list Match entries of prefix-lists
```

upvoted 1 times

Stivostine 2 years ago

D is ok.

It's written for ex : match ip adress WAN_PREFIXES
and should be : match ip adress prefix-list WAN_PREFIXES

Same for LOCAL_PREFIXES & VPN_PREFIXES

upvoted 3 times

JOKERR 2 years ago

D is corrrect. Prefix-list in route-maps shoujld be specified using prefix list keyword. Otherwise, route-map takes it for access-list

```
ER1(config-route-map)#match ip address ?
<1-199> IP access-list number
<1300-2699> IP access-list number (expanded range)
WORD IP access-list name
prefix-list Match entries of prefix-lists
```

upvoted 2 times

Question #37 *Topic 1*

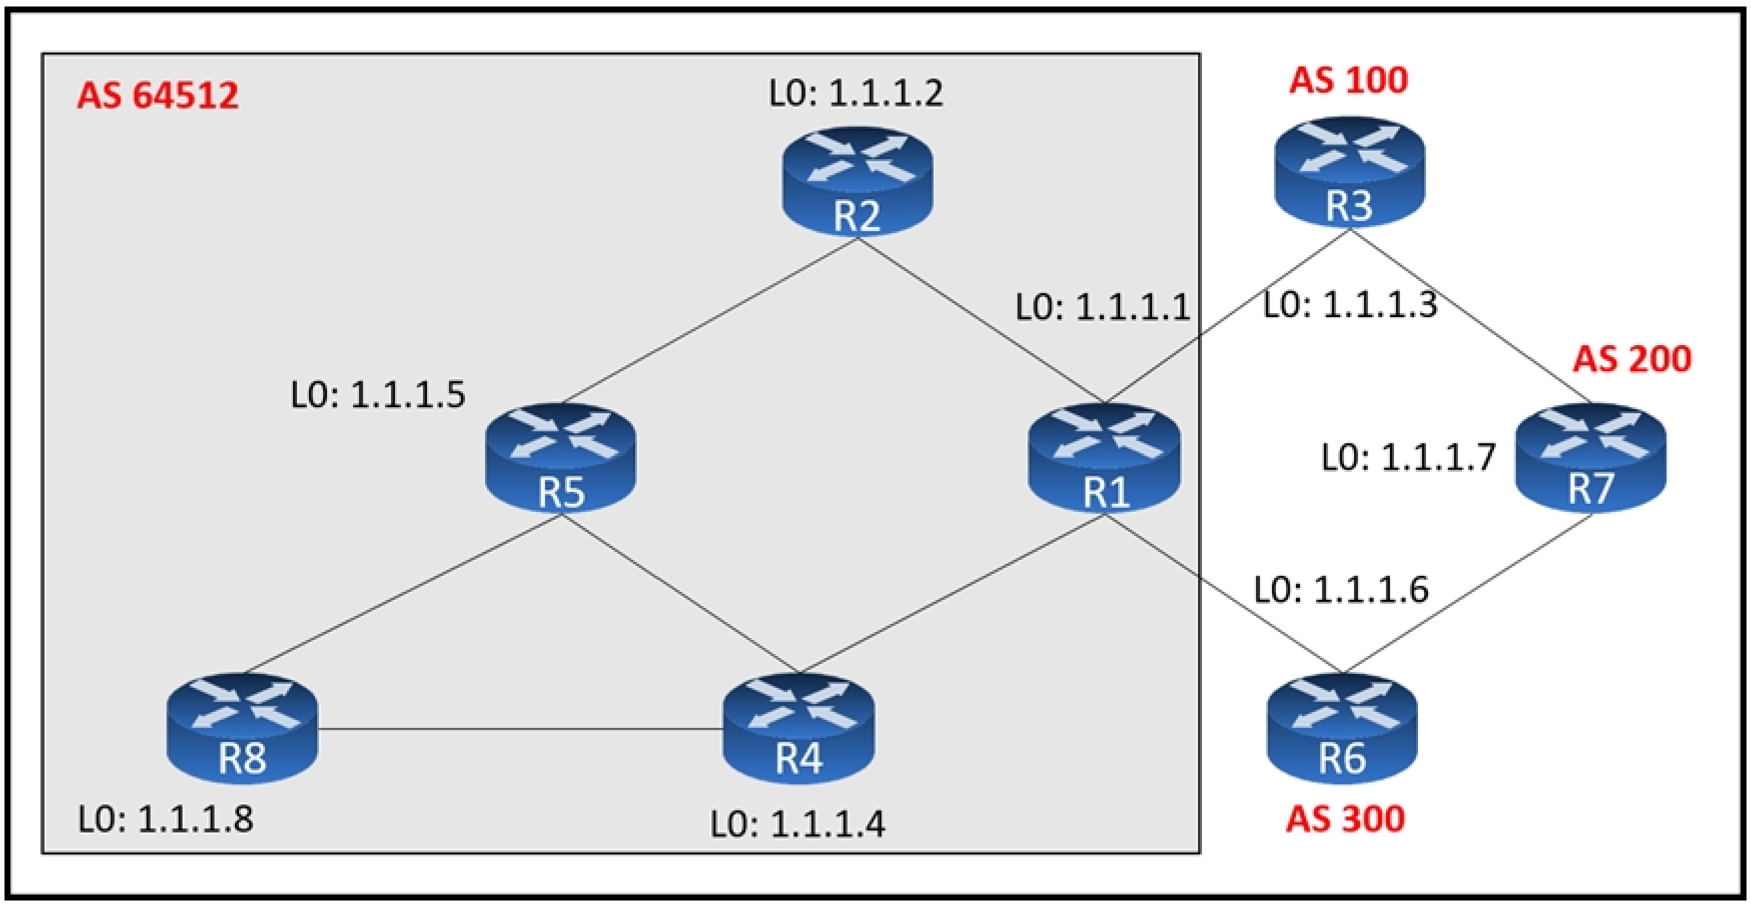

Refer to the exhibit. An engineer configured R2 and R5 as route reflectors and noticed that not all routes are sent to R1 to advertise to the eBGP peers.
Which iBGP routers must be configured as route reflectors to advertise all routes to restore reachability across all networks?

A. R1 and R4

B. R1 and R5

C. R4 and R5

D. R2 and R5

Correct Answer: *C*

Community vote distribution

C (67%) A (33%)

studybuddy10 Highly Voted 2 years, 1 month ago

C - confirmed in the lab.

All that is needed is R4 to be a RR with R1 as its client and that gets all loopbacks in routers BGP tables. So having R2, R5 and R4 also works. R4 and R5 is the only option that works without any other RR configuration. So the answers assume we roll back the engineers config and take a fresh start. Definitely C.

upvoted 6 times

studybuddy10 2 years, 1 month ago

Tested further, only R4 is needed as RR as a minimum. With what Alimo123 says below, its bad practice to have an edge router as RR so that would eliminate answer A, still C as the answer, but R5 is not needed as RR.

upvoted 4 times

gndrx78 1 year, 11 months ago

Apart from the test, what is the reason? RR rules do not help here apart from the fact the two RRs must be in a mesh, that means connected (so it cannot be R1 and R5). If we exclude answers B and D we have solutions A and C that means R4 and R1 or R5. We could exclude R1 because it is better not to use an edge router as RR but I have not found any real reason to choose C in accordance with RR rules.
Explanation found here:
https://itexamanswers.net/ccnp-enarsi-300-410-dumps-full-questions-with-vce-pdf.html/2
seems to be wrong because RR have to speak to each other. So far, I cannot really say answer is C

upvoted 1 times

HungarianDish Most Recent 7 months ago

Selected Answer: C

I also confirmed solution "C" in the lab (CML). RR = R4 (it's clients = R1, R5, R8) and RR = R5 (it's client R2). R2's loopback has not been advertised to R4 and R8 (and vica versa) until R2 became the client of RR R5. In my lab, full reachability was achieved with R4 and R5 being RRs. (Maybe I am missing something as others stated that R4 would be enough as RR.)

upvoted 4 times

AinsB 7 months, 2 weeks ago

Selected Answer: A

There is no difference between the setup of R2 & R5 and R4 & R5. So the same problem would exist with C. R1 R4 (A) would be better but it would still be a poor design. Based on the diagram the best answer would be R2 & R4.

upvoted 2 times

toto89 11 months, 2 weeks ago

I think the answer is C too. If R5 and R2 have the same cluster-id, then R8 loopback coming from R5 will never be advertised to R1 because it will be discarded by R2.
The real interesting question here is WHY do we have a question like this ? Multiple route reflectors topic is excluded in the exam topics.. And why don't they say that the cluster-id is the same ?
Poor cisco question as always. Makes me want to dump more.

upvoted 1 times

gndrx78 1 year, 6 months ago

Hello, I made some tests with GNS3 and it seems A,C and D are good solutions since all loopback interfaces are reachable. But answer D is excluded by the exercise and R1 is better not to use as RR. So the only answer remaining is C. Studybuddy is correct when he says only R4 is enough but if you make R1 as client, R2 is not reachable. The right thing to do is set R4 as RR and set R5 and R8 as RR client. I hope it may help.

upvoted 1 times

timtgh 1 year, 6 months ago

Just R4 as RR with R1, R5, and R8 as clients would work, but is not an option. Options C and D both work, but C seems more likely to be what they want.

upvoted 2 times

gndrx78 1 year, 6 months ago

Hi timtgh, R1 is not necessary as client since it is already connected to the rest of the network.

upvoted 1 times

AliMo123 1 year, 10 months ago

```
R2:
interface FastEthernet0/0
ip address 192.168.2.2 255.255.255.0
duplex half
!
interface FastEthernet1/0
ip address 192.168.5.2 255.255.255.0
duplex half
!
router bgp 400
bgp log-neighbor-changes
network 192.168.2.0
network 192.168.5.0
neighbor 192.168.2.1 remote-as 400
neighbor 192.168.5.5 remote-as 400
neighbor 192.168.5.5 route-reflector-client

R4:
interface FastEthernet0/0
ip address 192.168.1.4 255.255.255.0
duplex half
!
interface FastEthernet1/0
ip address 192.168.4.4 255.255.255.0
duplex half
!
interface FastEthernet2/0
ip address 192.168.10.4 255.255.255.0
duplex half
!
router bgp 400
bgp log-neighbor-changes
network 192.168.1.0
network 192.168.4.0
network 192.168.10.0
neighbor 192.168.1.1 remote-as 400
neighbor 192.168.4.8 remote-as 400
neighbor 192.168.4.8 route-reflector-client
neighbor 192.168.10.5 remote-as 400
```

R2 &R4 are working perfectly fine

upvoted 1 times

wts 1 year, 11 months ago

Why can't RR2 send routes received from RR5 to R1?

upvoted 1 times

wts 1 year, 11 months ago

Why not R2 and R4?
upvoted 1 times

AliMo123 1 year, 10 months ago
I did lab and R2 and R4 work perfectly fine
upvoted 2 times

wts 1 year, 9 months ago
Then I would replace in the question "routers must be" with "routers can be" or somehow change it.
upvoted 1 times

testbench007 1 year, 11 months ago
A poorly worded and considered example. but i would settle for C. R4 has to be a RR
upvoted 2 times

markan 2 years ago
I really dont understand why R2 and R5 as RR don't work.
upvoted 4 times

timtgh 1 year, 6 months ago
That would work also. But probably not the answer they are looking for.
upvoted 1 times

AliMo123 2 years, 1 month ago
It is not a good practice to have an edge router as RR
upvoted 3 times

Raider1 2 years, 1 month ago
It more sense configure reflectors on R1 and R5
upvoted 2 times

JOKERR 2 years ago
No it does not. Because R1 and R5 cannot exchange routes since they are iBGP neighbors.
upvoted 2 times

gndrx78 1 year, 11 months ago
JOKERR is right:
It is important to note that route reflectors must form a full mesh connectivity among themselves and each client peer with only its route reflector. Full mesh among route reflectors is not apparent until there are at least three route reflectors (see Figure 9.12(d)).
Source: https://www.sciencedirect.com/topics/computer-science/route-reflector
upvoted 1 times

wts 1 year, 11 months ago
The client does not know that he is a client, he simply sends routes towards the neighbor.
upvoted 1 times

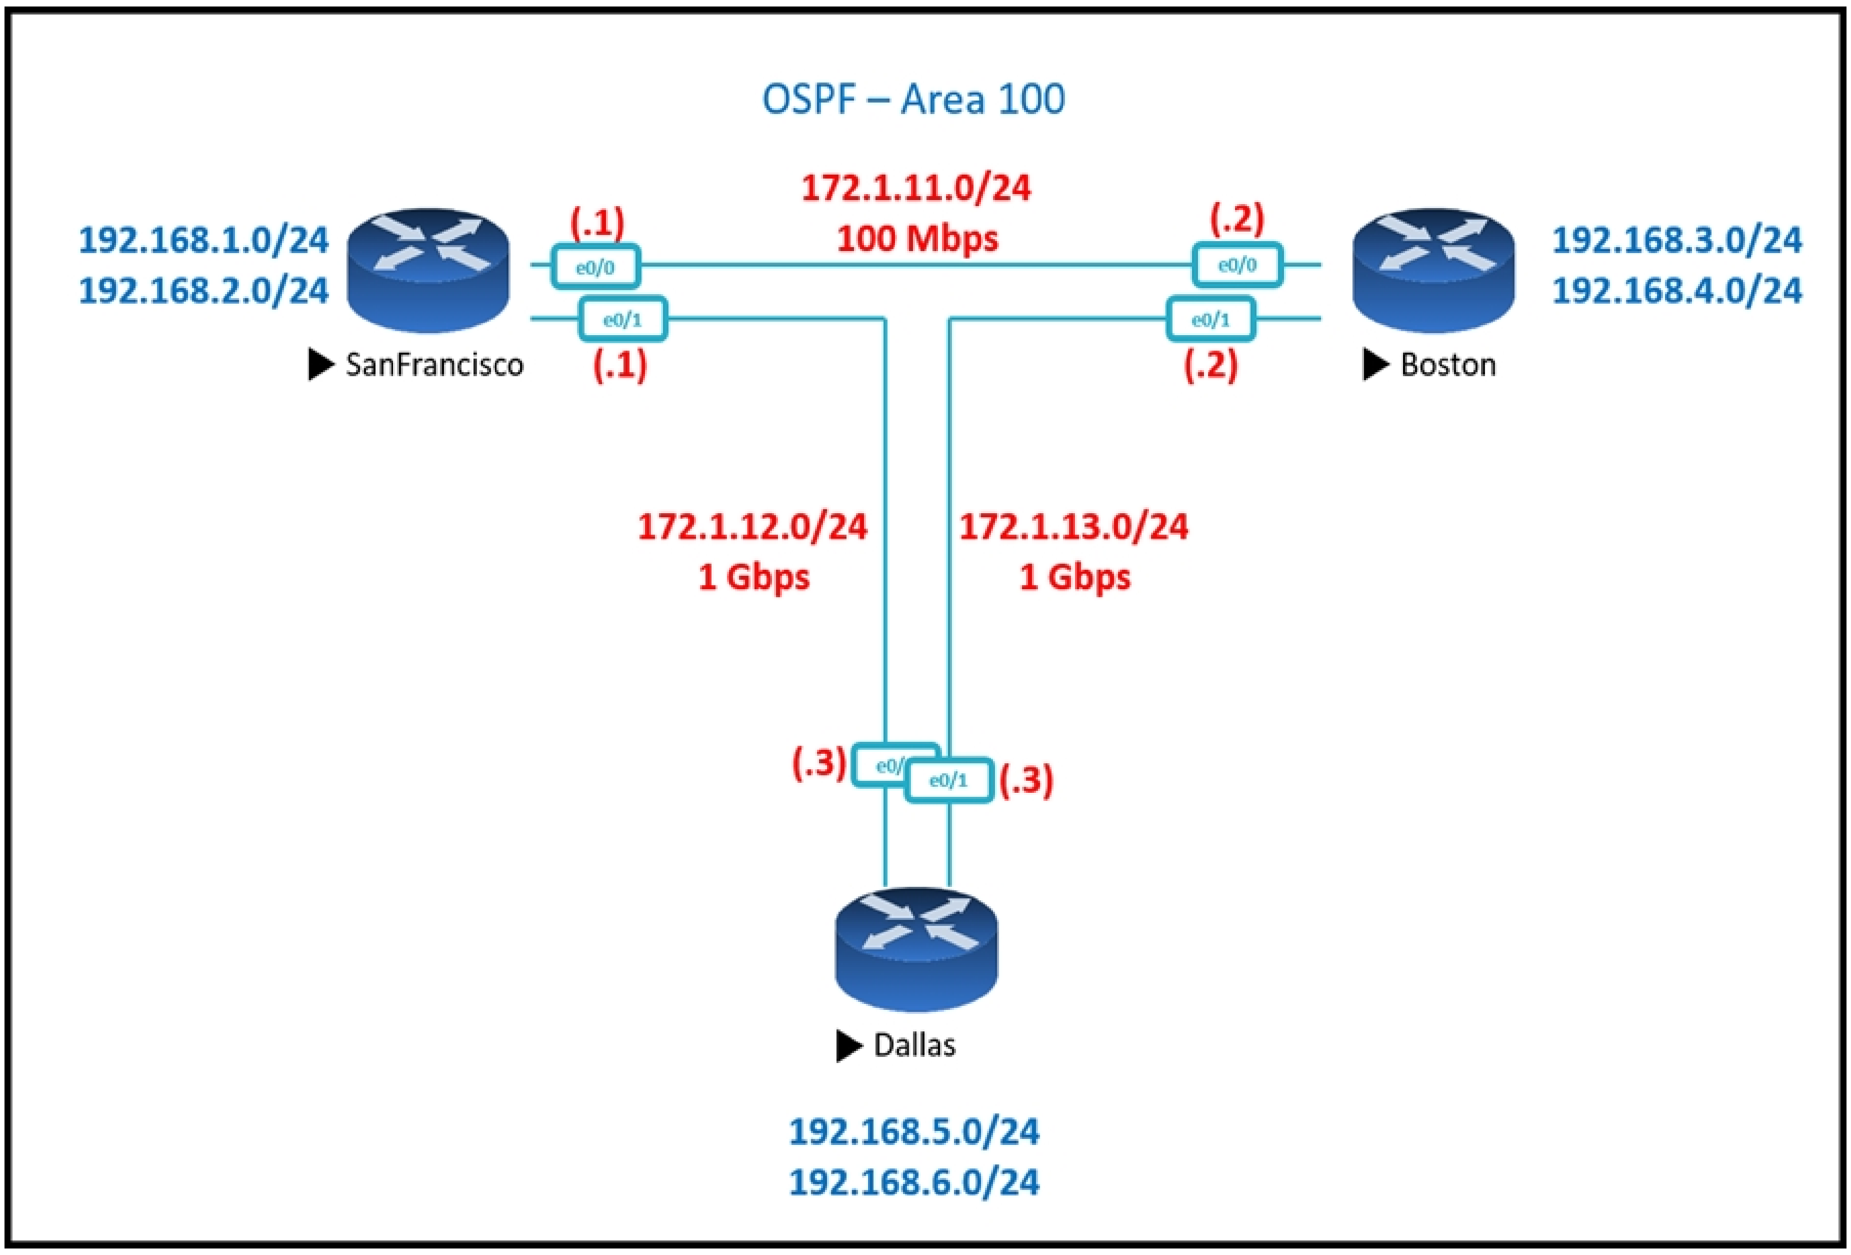
OSPF – Area 100
172.1.11.0/24
100 Mbps
192.168.1.0/24
192.168.2.0/24
(.1)
e0/0
e0/1
(.1)
SanFrancisco
(.2)
e0/0
e0/1
(.2)
Boston
192.168.3.0/24
192.168.4.0/24
172.1.12.0/24
1 Gbps
172.1.13.0/24
1 Gbps
(.3)
e0/
e0/1
(.3)
Dallas
192.168.5.0/24
192.168.6.0/24

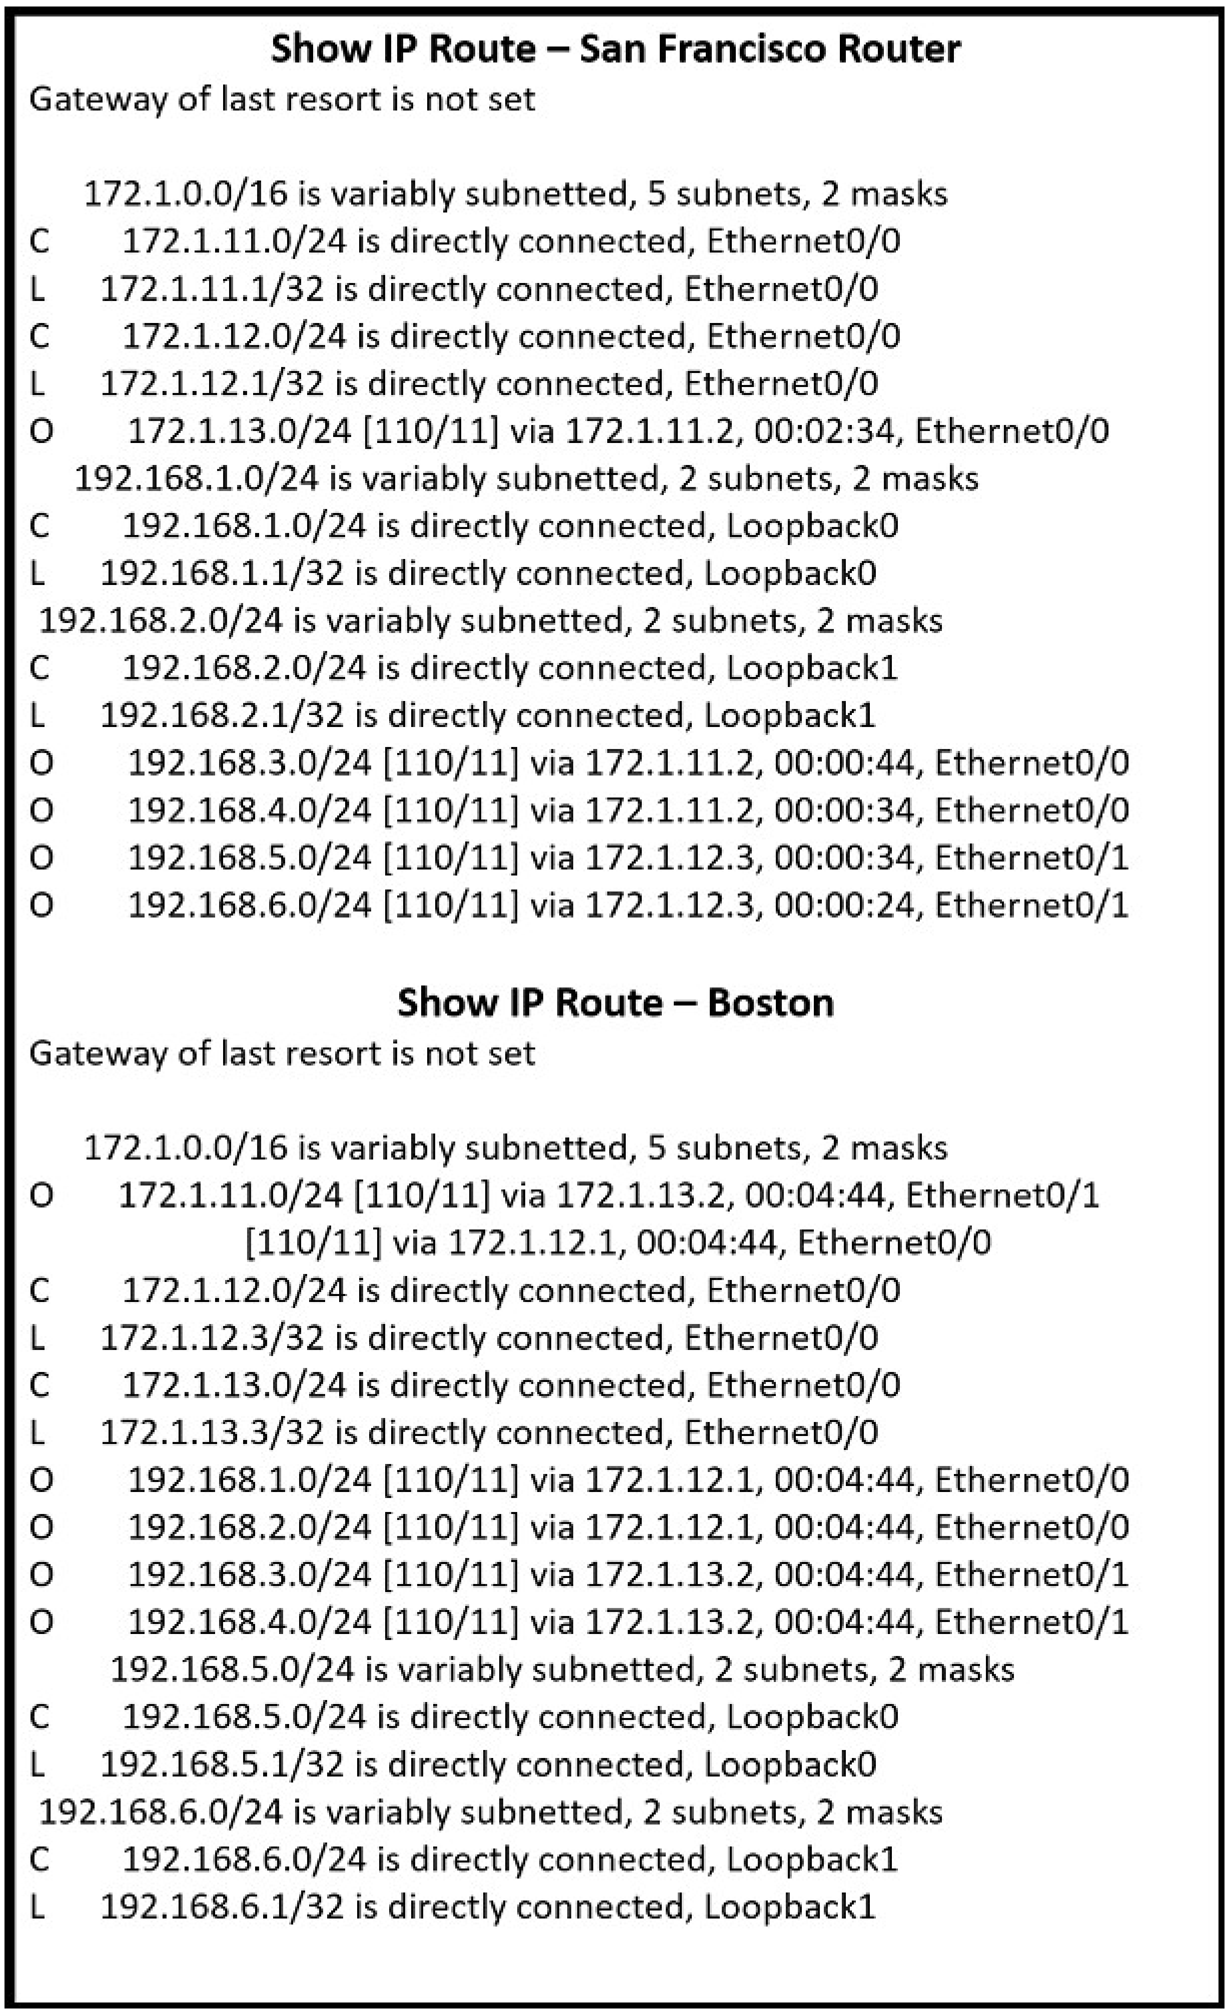

```
Show IP Route – San Francisco Router
Gateway of last resort is not set

     172.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
C       172.1.11.0/24 is directly connected, Ethernet0/0
L     172.1.11.1/32 is directly connected, Ethernet0/0
C       172.1.12.0/24 is directly connected, Ethernet0/0
L     172.1.12.1/32 is directly connected, Ethernet0/0
O       172.1.13.0/24 [110/11] via 172.1.11.2, 00:02:34, Ethernet0/0
    192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.1.0/24 is directly connected, Loopback0
L     192.168.1.1/32 is directly connected, Loopback0
 192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.2.0/24 is directly connected, Loopback1
L     192.168.2.1/32 is directly connected, Loopback1
O       192.168.3.0/24 [110/11] via 172.1.11.2, 00:00:44, Ethernet0/0
O       192.168.4.0/24 [110/11] via 172.1.11.2, 00:00:34, Ethernet0/0
O       192.168.5.0/24 [110/11] via 172.1.12.3, 00:00:34, Ethernet0/1
O       192.168.6.0/24 [110/11] via 172.1.12.3, 00:00:24, Ethernet0/1

Show IP Route – Boston
Gateway of last resort is not set

     172.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
O      172.1.11.0/24 [110/11] via 172.1.13.2, 00:04:44, Ethernet0/1
                 [110/11] via 172.1.12.1, 00:04:44, Ethernet0/0
C       172.1.12.0/24 is directly connected, Ethernet0/0
L     172.1.12.3/32 is directly connected, Ethernet0/0
C       172.1.13.0/24 is directly connected, Ethernet0/0
L     172.1.13.3/32 is directly connected, Ethernet0/0
O       192.168.1.0/24 [110/11] via 172.1.12.1, 00:04:44, Ethernet0/0
O       192.168.2.0/24 [110/11] via 172.1.12.1, 00:04:44, Ethernet0/0
O       192.168.3.0/24 [110/11] via 172.1.13.2, 00:04:44, Ethernet0/1
O       192.168.4.0/24 [110/11] via 172.1.13.2, 00:04:44, Ethernet0/1
       192.168.5.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.5.0/24 is directly connected, Loopback0
L     192.168.5.1/32 is directly connected, Loopback0
 192.168.6.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.6.0/24 is directly connected, Loopback1
L     192.168.6.1/32 is directly connected, Loopback1
```

Refer to the exhibits. SanFrancisco and Boston routers are choosing slower links to reach each other despite the direct links being up. Which configuration fixes the issue?

A. All Routers router ospf 1 auto-cost reference-bandwidth 100

B. SanFrancisco Router router ospf 1 auto-cost reference-bandwidth 1000

C. Boston Router router ospf 1 auto-cost reference-bandwidth 1000

D. All Routers router ospf 1 auto-cost reference-bandwidth 1000

Correct Answer: *D*

Community vote distribution

D (100%)

Stivostine Highly Voted 2 years ago

D is ok :

Under the OSPF process, the command auto-cost reference-bandwidth bandwidth-in-mbps changes the reference bandwidth for all OSPF interfaces associated with that process. If the reference bandwidth is changed on one router, then the reference bandwidth should be

changed on all OSPF routers to ensure that SPF uses the same logic to prevent routing loops.
It is a best practice to set the same reference bandwidth for all OSPF routers.

upvoted 5 times

Brand Most Recent 3 months, 4 weeks ago

Selected Answer: D

```
R1(config-router)#auto-cost reference-bandwidth ?
<1-4294967> The reference bandwidth in terms of Mbits per second
R1(config-router)#auto-cost reference-bandwidth
```

it's being defined as mbps, so it's D

upvoted 1 times

Noproblem22 1 year, 1 month ago

They have made a mistake, the output of the lower router is Dallas. D is correct

upvoted 2 times

yuki0829 1 year, 4 months ago

I'think the lower route table is not Boston's.
It's Dallas's.

upvoted 2 times

timtgh 1 year, 6 months ago

Poorly worded. They mean even though the faster links are up, not "direct" links. The direct links are the slower links. Anyway, D is right, because on many Cisco router platforms, the reference bandwidth default is less than 1000, which makes it inaccurate for 1Gb links or higher.

upvoted 3 times

Question #39 *Topic 1*

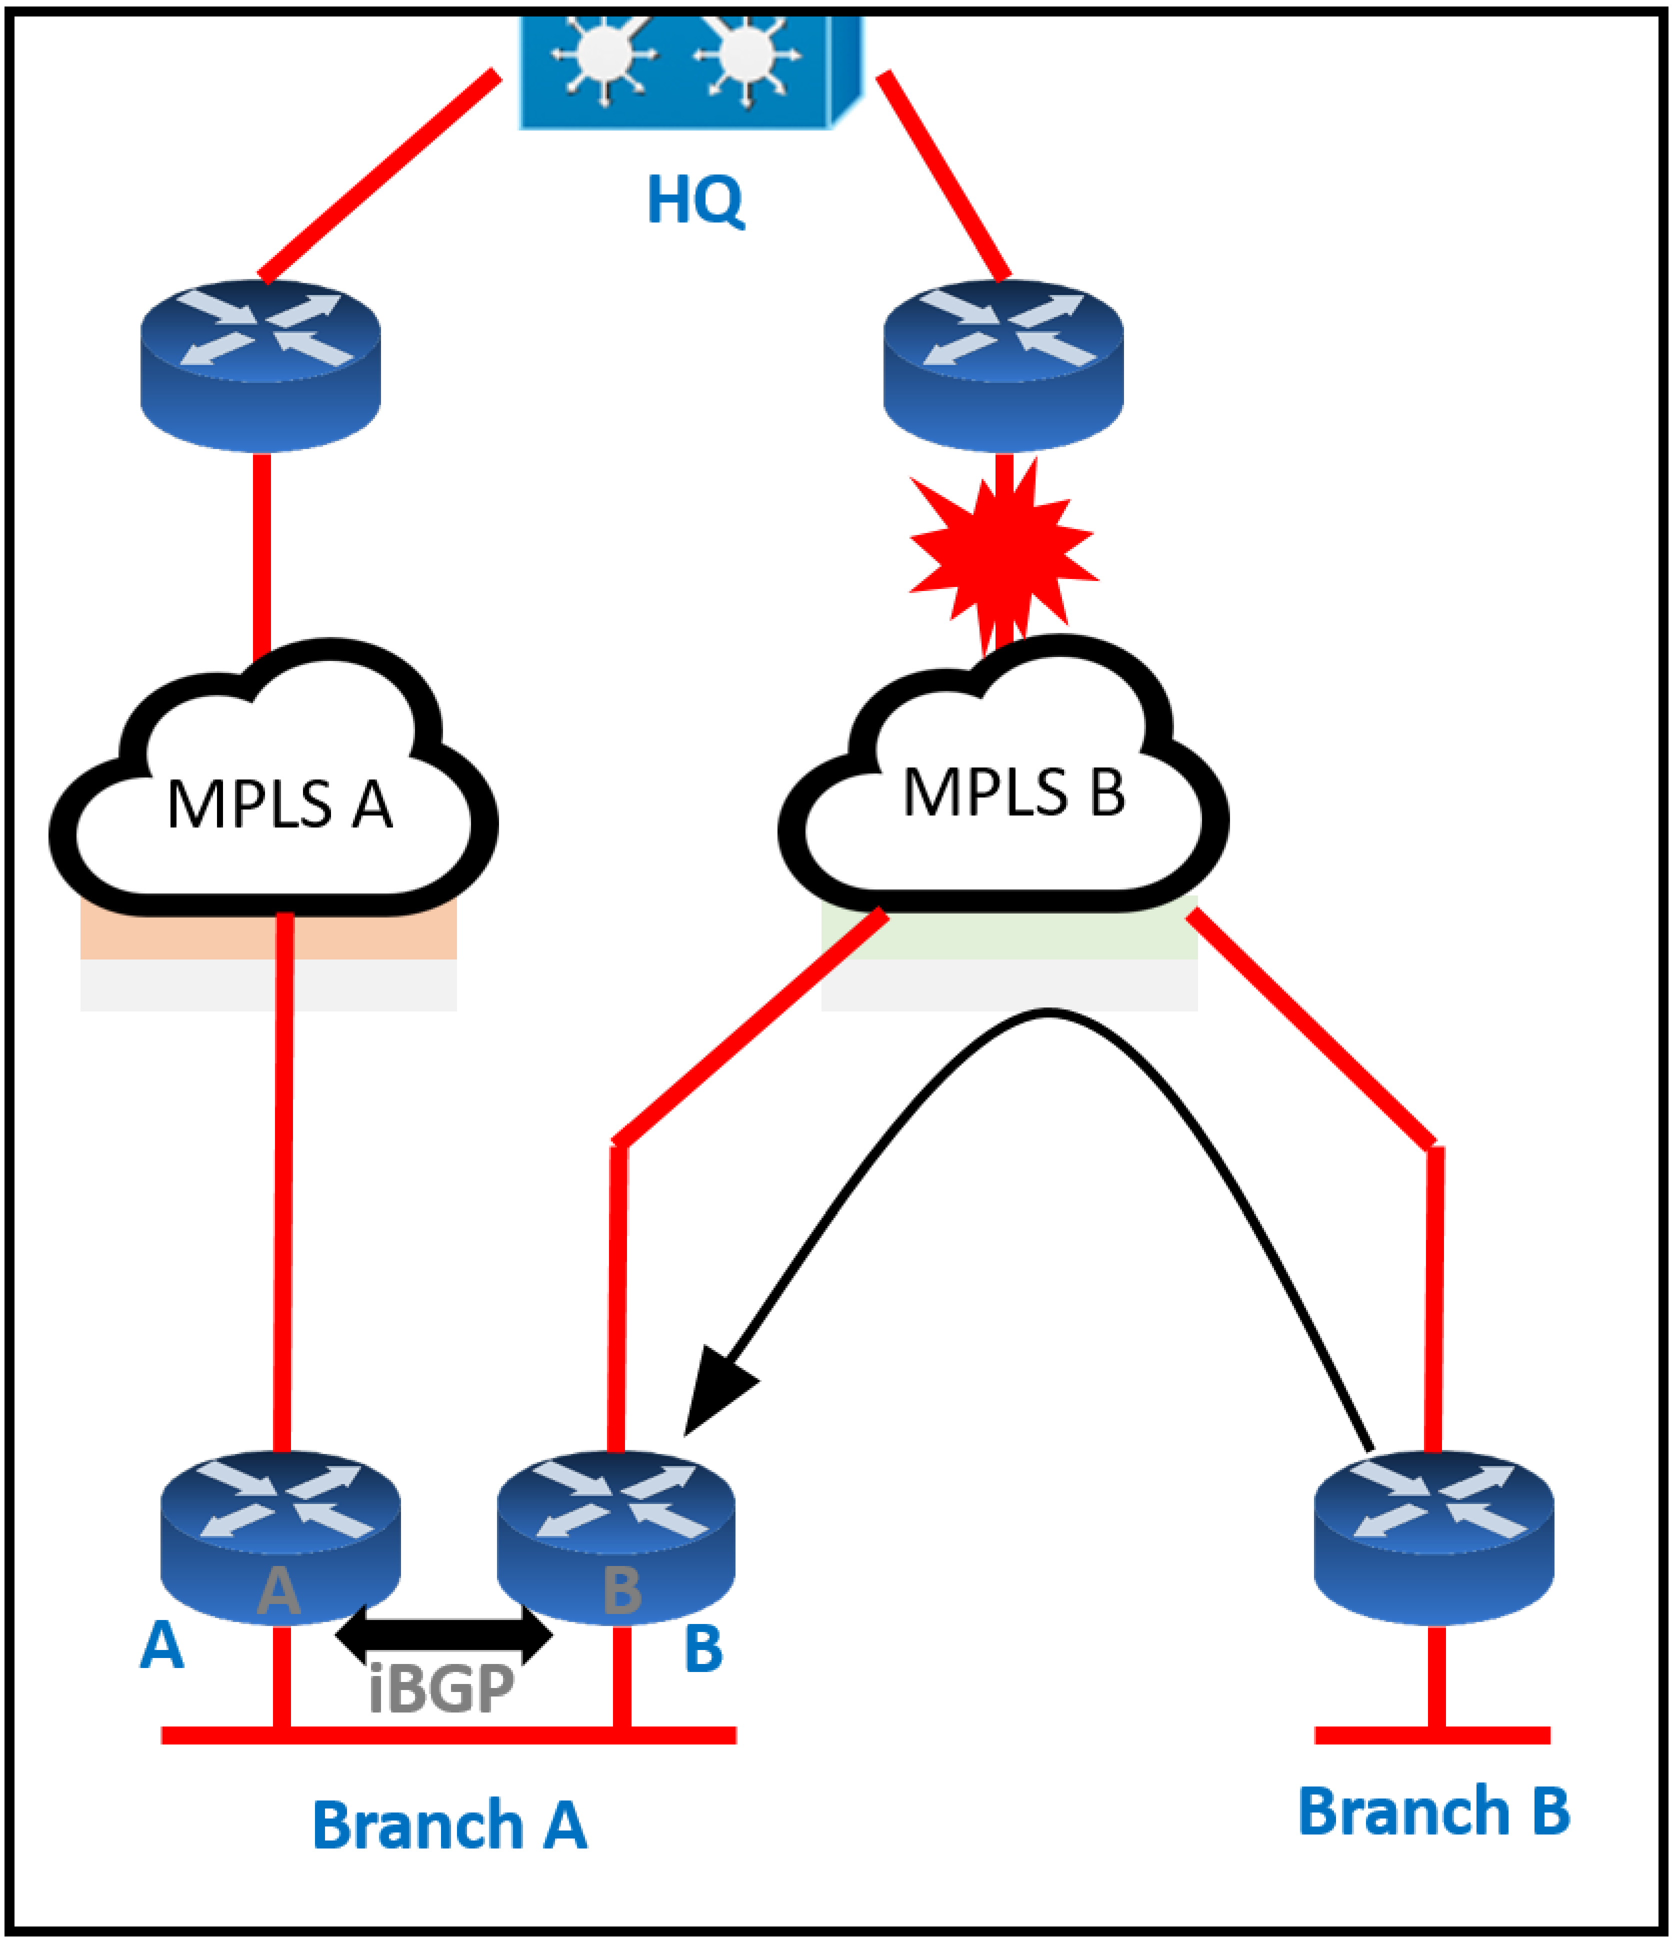

Refer to the exhibit. Troubleshoot and ensure that branch I’ only ever uses the MPLS I’ network to reach HQ.
Which action achieves this requirement?

A. Introduce AS path prepending on the branch A MPLS I’ network connection so that any HQ advertisements from branch A toward the MPLS I’ network are prepended three times

B. Modify the weight of all HQ prefixes received at branch I’ from the MPLS I’ network to be higher than the weights used on the MPLS A network

C. Increase the local preference for all HQ prefixes received at branch I’ from the MPLS I’ network to be higher than the local preferences used on the MPLS A network

D. Introduce an AS path filter on branch A routers so that only local prefixes are advertised into BGP

Correct Answer: *B*

Community vote distribution

D (100%)

⊟ Alnet Highly Voted 2 years, 1 month ago

Answer D seems most logical. Question says that Branch B ONLY EVER uses MPLS B. That means you don't want the path through Branch A as an alternate. So A, B and C all prefer HQ routes, but they don't eliminate Branch A routes. D is the only answer which actually filters the route.

upvoted 9 times

⊟ **JOKERR** 2 years ago

Yes. D seems correct because the question says: "only ever uses". So I think it means it should not use MPLS A at all.

upvoted 2 times

⊟ **[Removed]** 2 years ago

Yea I cant see it being any answer other than D. Completely filter out the AS and only advertising the local prefixes ensures that through MPLS B will be on the only option in the bgp table...

upvoted 1 times

⊟ **XBfoundX** Most Recent ⏲ 6 months ago

as many users as said the only answer that can be the true one is D the other answers are just make the route less prefered, by the way we assume that the branch is using another ASN number, because if the ASN is the same the advertisement will be sent anyway because they are ibgp neighbors and when you receive and ibgp update the ASN number is not specified in the update (remember that if also Branch B will be an ibgp neighbor to configure RR in RB)

upvoted 1 times

⊟ **anonymous1966** 8 months, 3 weeks ago

For me the site codification has a problem.
The text with correct codification is below:
Troubleshoot and ensure that branch B only ever uses the MPLS B network to reach HQ.

Which action achieves this requirement?

(A) Modify the weight of all HQ prefixes received at branch B from the MPLS B network to be higher than the weights used on the MPLS A network
(B) Increase the local preference for all HQ prefixes received at branch B from the MPLS B network to be higher than the local preferences used on the MPLS A network
(C) Introduce AS path prepending on the branch A MPLS B network connection so that any HQ advertisements from branch A toward the MPLS B network are prepended three times
(D) Introduce an AS path filter on branch A routers so that only local prefixes are advertised into BGP

upvoted 1 times

⊟ **Koume** 11 months, 1 week ago

Selected Answer: D

The only way to be shure that branch B do do not use Branch A as transit is jus filter Branch A to filter the annoumcent to HQ, So D is correct

upvoted 1 times

⊟ **leogp79** 1 year, 4 months ago

Branch A with option D, is not transit AS for the other two AS, thus BRANCH B always goes to MPLS B to HQ

upvoted 1 times

⊟ **Edwinmolinab** 1 year, 5 months ago

Selected Answer: D

To me B is not correct because in this case we're using a MPLS network and the weight variable only exits on cisco equipment and is for local use, and D option seemed most appropiate.

upvoted 1 times

⊟ **thanh123** 1 year, 8 months ago

After reading the questions many times, D is seem to be the correct one. Other answers just make the route via MPLS B is prefer to MPLS A. If the link is down, B will go to A to HQ. So D is the correct one

upvoted 1 times

⊟ **Carl1999** 1 year, 10 months ago

Selected Answer: D

B is wrong. route from branch A goes to MPLS B.
D is correct.

upvoted 1 times

⊟ **Hack4** 1 year, 10 months ago

B is the correct one

upvoted 1 times

⊟ **wts** 1 year, 10 months ago

Selected Answer: D

A - the path through branch A is worsened by the lengthening of AS-path, keeping the alternative. Condition "only ever uses" is not met.
B - when the route with the best WEIGHT disappears from the table, traffic will flow through branch A. Condition "only ever uses" is not met.
C - the path through branch A is still preserved, the path through MPLS B is being improved by LP.
D(correct answer) - office A advertises up and to the right of the picture only its own routes, branch B does not have an alternative route to HQ, branch B ONLY EVER USES the MPLS B network to reach HQ.

* - I assume that I' is B.

upvoted 3 times

⊟

Networkingguy 1 year, 10 months ago

Selected Answer: D

Alnet is correct with his explanation, D is correct it would seem

upvoted 1 times

Jenia1 1 year, 10 months ago

B seems to be the correct one
According to Image, you should be routed via Branch A
D is incorrect, if you Introduce an AS path filter on branch A routers so that only local prefixes are advertised into BGP, how the branch B will know the path via Branch A to the Network HQ?

upvoted 1 times

[Removed] 1 year, 10 months ago

Refer to the exhibit. Troubleshoot and ensure that branch B only ever uses the MPLS B network to reach HQ. Which action achieves this requirement?

Thats the entire question.. Only D will satisfy this. The other answers just make the route less preferable oppose to completely removing.

upvoted 1 times

Jenia1 1 year, 10 months ago

I see, thanks, looks like this image is broken, makes think it should go via branch A, and not directly to HQ as we can see the red sign.

upvoted 1 times

Tuchi 2 years ago

B is the corect one.

upvoted 1 times

branbush 2 years, 1 month ago

Isn't Answer D correct?

upvoted 1 times

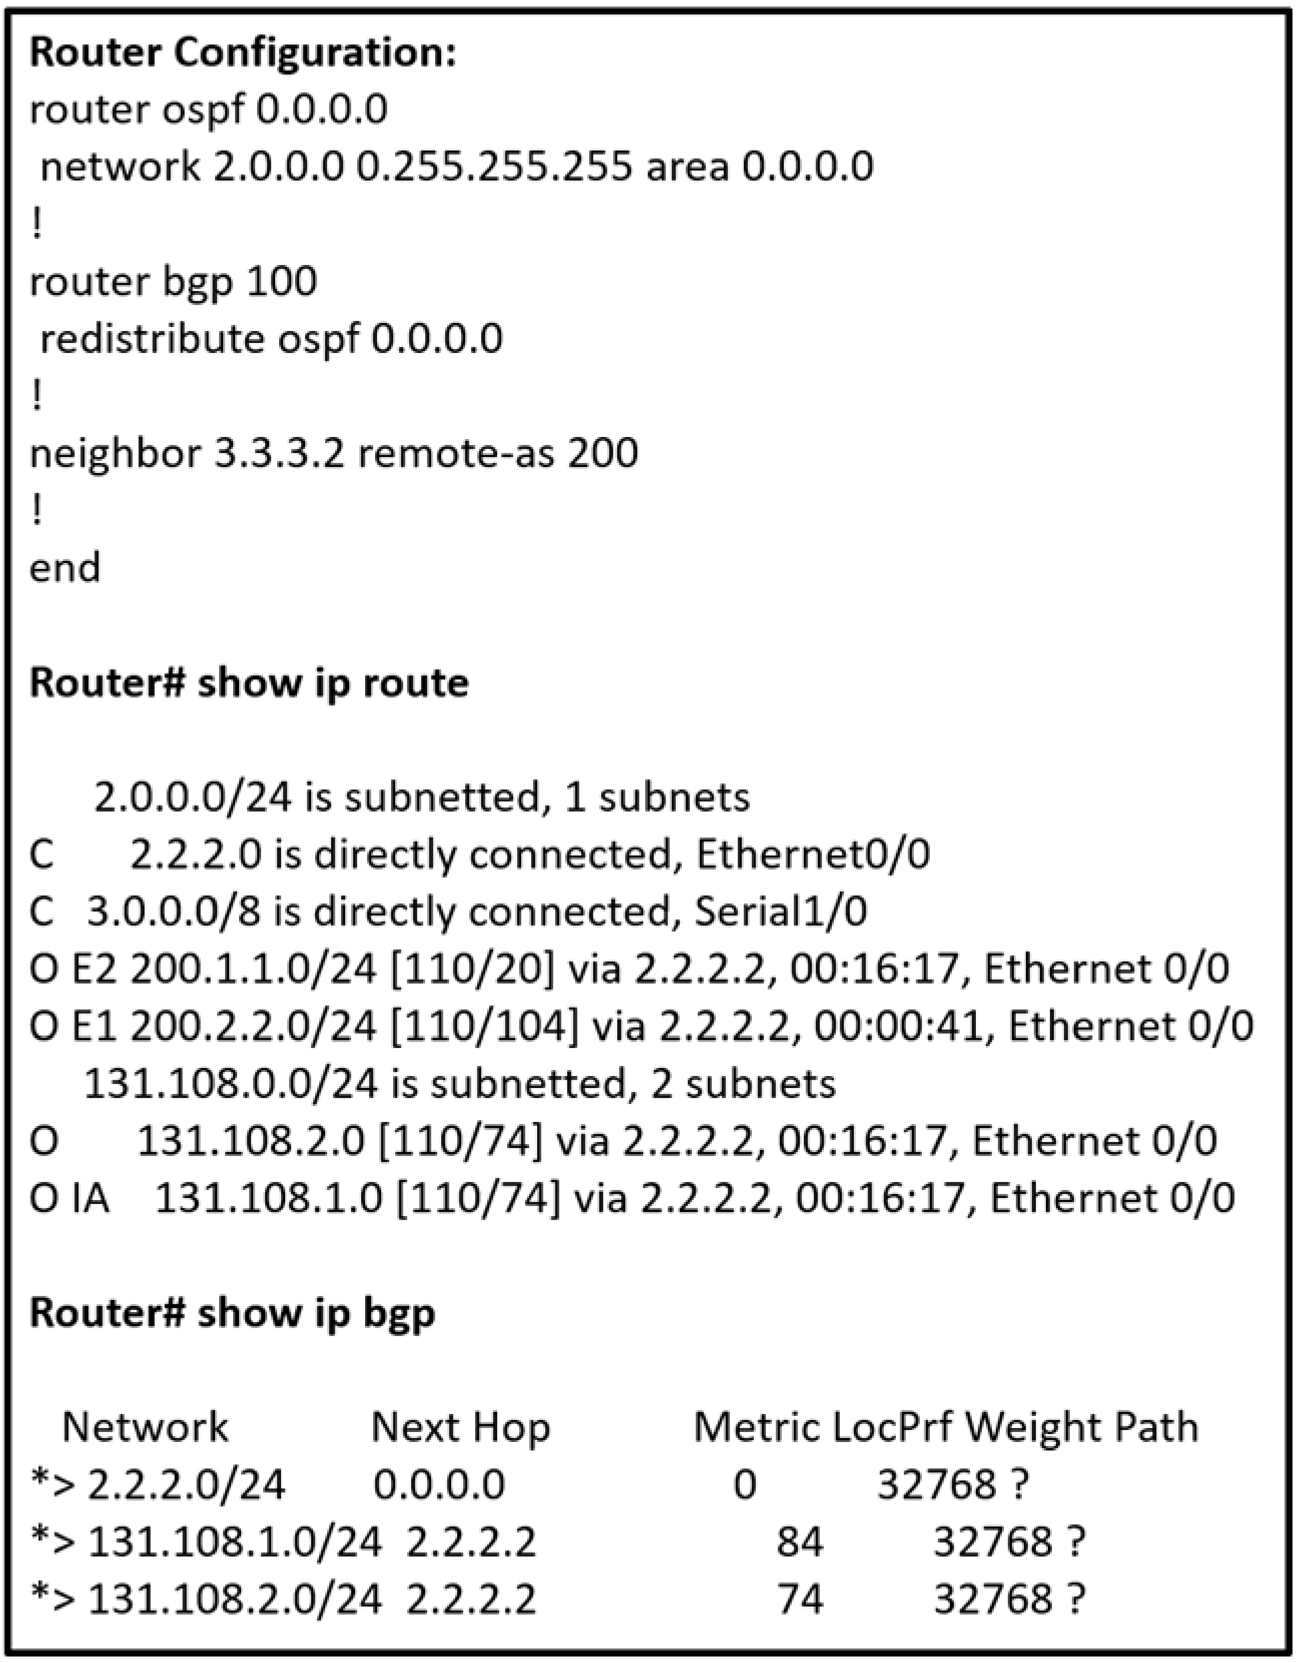

Router Configuration:

```
router ospf 0.0.0.0
 network 2.0.0.0 0.255.255.255 area 0.0.0.0
!
router bgp 100
 redistribute ospf 0.0.0.0
!
neighbor 3.3.3.2 remote-as 200
!
end
```

Router# show ip route

```
     2.0.0.0/24 is subnetted, 1 subnets
C       2.2.2.0 is directly connected, Ethernet0/0
C    3.0.0.0/8 is directly connected, Serial1/0
O E2 200.1.1.0/24 [110/20] via 2.2.2.2, 00:16:17, Ethernet 0/0
O E1 200.2.2.0/24 [110/104] via 2.2.2.2, 00:00:41, Ethernet 0/0
     131.108.0.0/24 is subnetted, 2 subnets
O       131.108.2.0 [110/74] via 2.2.2.2, 00:16:17, Ethernet 0/0
O IA    131.108.1.0 [110/74] via 2.2.2.2, 00:16:17, Ethernet 0/0
```

Router# show ip bgp

```
   Network          Next Hop          Metric LocPrf Weight Path
*> 2.2.2.0/24       0.0.0.0              0          32768 ?
*> 131.108.1.0/24   2.2.2.2                84          32768 ?
*> 131.108.2.0/24   2.2.2.2                74          32768 ?
```

Refer to the exhibit. The OSPF routing protocol is redistributed into the BGP routing protocol, but not all the OSPF routes are distributed into BGP.

Which action resolves the issue?

A. Include the word external in the redistribute command

B. Use a route-map command to redistribute OSPF external routes defined in an access list

C. Include the word internal external in the redistribute command

D. Use a route-map command to redistribute OSPF external routes defined in a prefix list

Correct Answer: *C*

Community vote distribution

C (65%) A (35%)

MrThinMints Highly Voted 1 year, 11 months ago

The material from cisco states first that: "If you configure the redistribution of OSPF into BGP without keywords, only OSPF intra-area and inter-area routes are redistributed into BGP, by default."
But then it says in order to distribute ONLY External Type 1 and Type 2 routes, you use the "external" keyword.
So reasoning on that, I am going with C. Include the word internal external in the redistribute command

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redis.html

upvoted 17 times

TECH3K3 1 year, 6 months ago

Shame you didn't lab it and find out the answer instead of being a bookworm.
I'm ashamed to say it, but this is the quality of the future cisco network engineers.
So many said the answer is C and you're all wrong and I've been a CCNP for years.
Every question if possible I try and lab and confirm the answer.
This si why we're going down the automation route because of LAZY network engineers.
we have free emulators from GNS3, eve-ng and lots of paid ones and we have people on here guessing and not bettering themselves by labbing.

upvoted 3 times

Jey117 2 months, 1 week ago

Hahaha.. A know-it-all who knows nothing.

upvoted 1 times

Brand 3 months, 4 weeks ago

```
R1(config-router-af)#redistribute ospf 1 match ?
external Redistribute OSPF external routes
internal Redistribute OSPF internal routes
nssa-external Redistribute OSPF NSSA external routes
R1(config-router-af)#redistribute ospf 1 match
```

you lazy CCNP...

upvoted 2 times

Almylle 5 months, 3 weeks ago

U are so wrong, i labbed it and the answer is C, if u want i can demostrate to u

upvoted 2 times

BTK0311 Most Recent 3 months ago

When redistributing OSPF routes into BGP, including the word "external" in the redistribute command typically resolves the issue when not all OSPF routes are being distributed into BGP. This is because the "external" keyword instructs BGP to redistribute OSPF external routes (routes from other autonomous systems) into BGP. If you omit "external," only OSPF internal routes (intra-area and inter-area routes) are redistributed by default.

Option B and Option D suggest using a route-map to control the redistribution of OSPF external routes based on specific criteria defined in an access list or prefix list. While these are valid methods to control redistribution, they do not directly address the issue of missing OSPF routes in BGP. Option C, "include the word internal external in the redistribute command," is not a standard syntax for redistribution and is not typically used in OSPF-to-BGP redistribution.

So, including the "external" keyword in the redistribute command is the most straightforward way to ensure that OSPF external routes are redistributed into BGP.

upvoted 1 times

Chiaretta 5 months ago

Selected Answer: C

C is correct

upvoted 1 times

MicMillon 5 months, 2 weeks ago

Selected Answer: C

C, you need internal and external. if you only specify external, it will only advertise external routes and you'll loose the internal ones

upvoted 2 times

MicMillon 5 months, 3 weeks ago

Selected Answer: C

C is correct

upvoted 1 times

cir_ 6 months, 3 weeks ago

Selected Answer: C

A will only redistribute external
C will redistribute internal & external

upvoted 1 times

Dacusai 7 months, 4 weeks ago

I just lab it and the thing is, If you run the command for the first time with the matching external key word only, it only redistribute the external routes. But if you use the redistribute ospf # with no keyword it will only pass the internal routes, and after doing this you use the command again with the match external key word only, you them will get the external also and it wont remove the internal ones. So if you want to run the command for the first time you need to use both, internal and external.

upvoted 1 times

upp3r 8 months, 1 week ago

All this confusion... just type in:

```
#router bgp 65500
#redistribute ospf 1 external
```

now view the output of "show ip protocols" and see ONLY external routes are redistributed

the answer is C
upvoted 2 times

Huntkey 1 year, 2 months ago
Both A and C are correct. Either with "external" only or with "internal external", both would be expanded to " redistribute ospf 1 match internal external 1 external 2". This is a bad question unless both A and C would be considered right
upvoted 1 times

jarz 1 year, 1 month ago
Only A is correct

Redistributing routes from OSPF to BGP does not include OSPF external routes by default match external [1 | 2] is required to redistribute OSPF external routes.
upvoted 1 times

Remsync 1 year, 2 months ago
The question says that all routes need to be redistributed, so C is the answer since A would only redistribute external routes and left out the internal ones.
upvoted 1 times

Remsync 1 year, 2 months ago
May bad. I labbed it, you only need to add external. A is also correct.
upvoted 2 times

doron1122 1 year, 4 months ago
c
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redis.html
upvoted 1 times

Deu_Inder 1 year, 4 months ago
Even I labbed it. The result is C.
There are some experienced CCNPs in this forum who say vehemently that it is A. They seem to have labbed it too. Just a small question: can the result be platform- and IOS dependent? I used C7200-ADVENTERPRISEK9-M under GNS3.
upvoted 1 times

Edwinmolinab 1 year, 5 months ago
If the command redistribute ospf is there and the administrator includes external the new line include internal and external 1 external 2, if the command doesn't exists when you apply the command only appears external 1 external 2 if the command already exists the new line only needs the external route for distribution. I probe it on GNS3
upvoted 1 times

TECH3K3 1 year, 5 months ago
lab it
upvoted 1 times

networkWiz 1 year, 5 months ago
Selected Answer: C
C. is the correct answer
even though you are no longer needed to add the "internal" keyword in latest versions (i think from v17).
For exam purpose ad both "internal" and "external" keyword in the command
upvoted 1 times

Reikidude00 1 year, 5 months ago
Selected Answer: A
Tested on GNS3, based on the same reference reported by MrThinMints, we can see that the documentation say

"Note: The configuration shows match external 1 external 2 and the command entered was redistribute ospf 1 match external. This is normal because OSPF automatically appends "external 1 external 2" in the configuration. It matches both OSPF external 1 and external 2 routes and it redistributes both routes into BGP. "

So A is correct
upvoted 1 times

Orchidium 1 year, 5 months ago
Selected Answer: A
Labbed this one and A is correct. You only need to add the "external" keyword. By default, internal and inter-area routes are redistributed from OSPF to BGP, so it is not necessary to add the internal keyword. Question 12 tackles the same issue.
upvoted 1 times

TECH3K3 1 year, 6 months ago
Selected Answer: A
The Answer is A [Include the word external in the redistribute command]
I have lab this and listen why.
You can NOT remove internal routes with the redistribution command even if you only specify external 1, external 2 or both.
You could remove the internal routes with a rout-map, but that's not what this question is asking.

WHY DON'T PEOPLE LAB TO CONFIRM AN ANSWER BEFORE SPREADING THEIR DUMB FAKE ANSWERS. You're supposed to be a Network Engineer, Lab Lab Lab..

upvoted 3 times

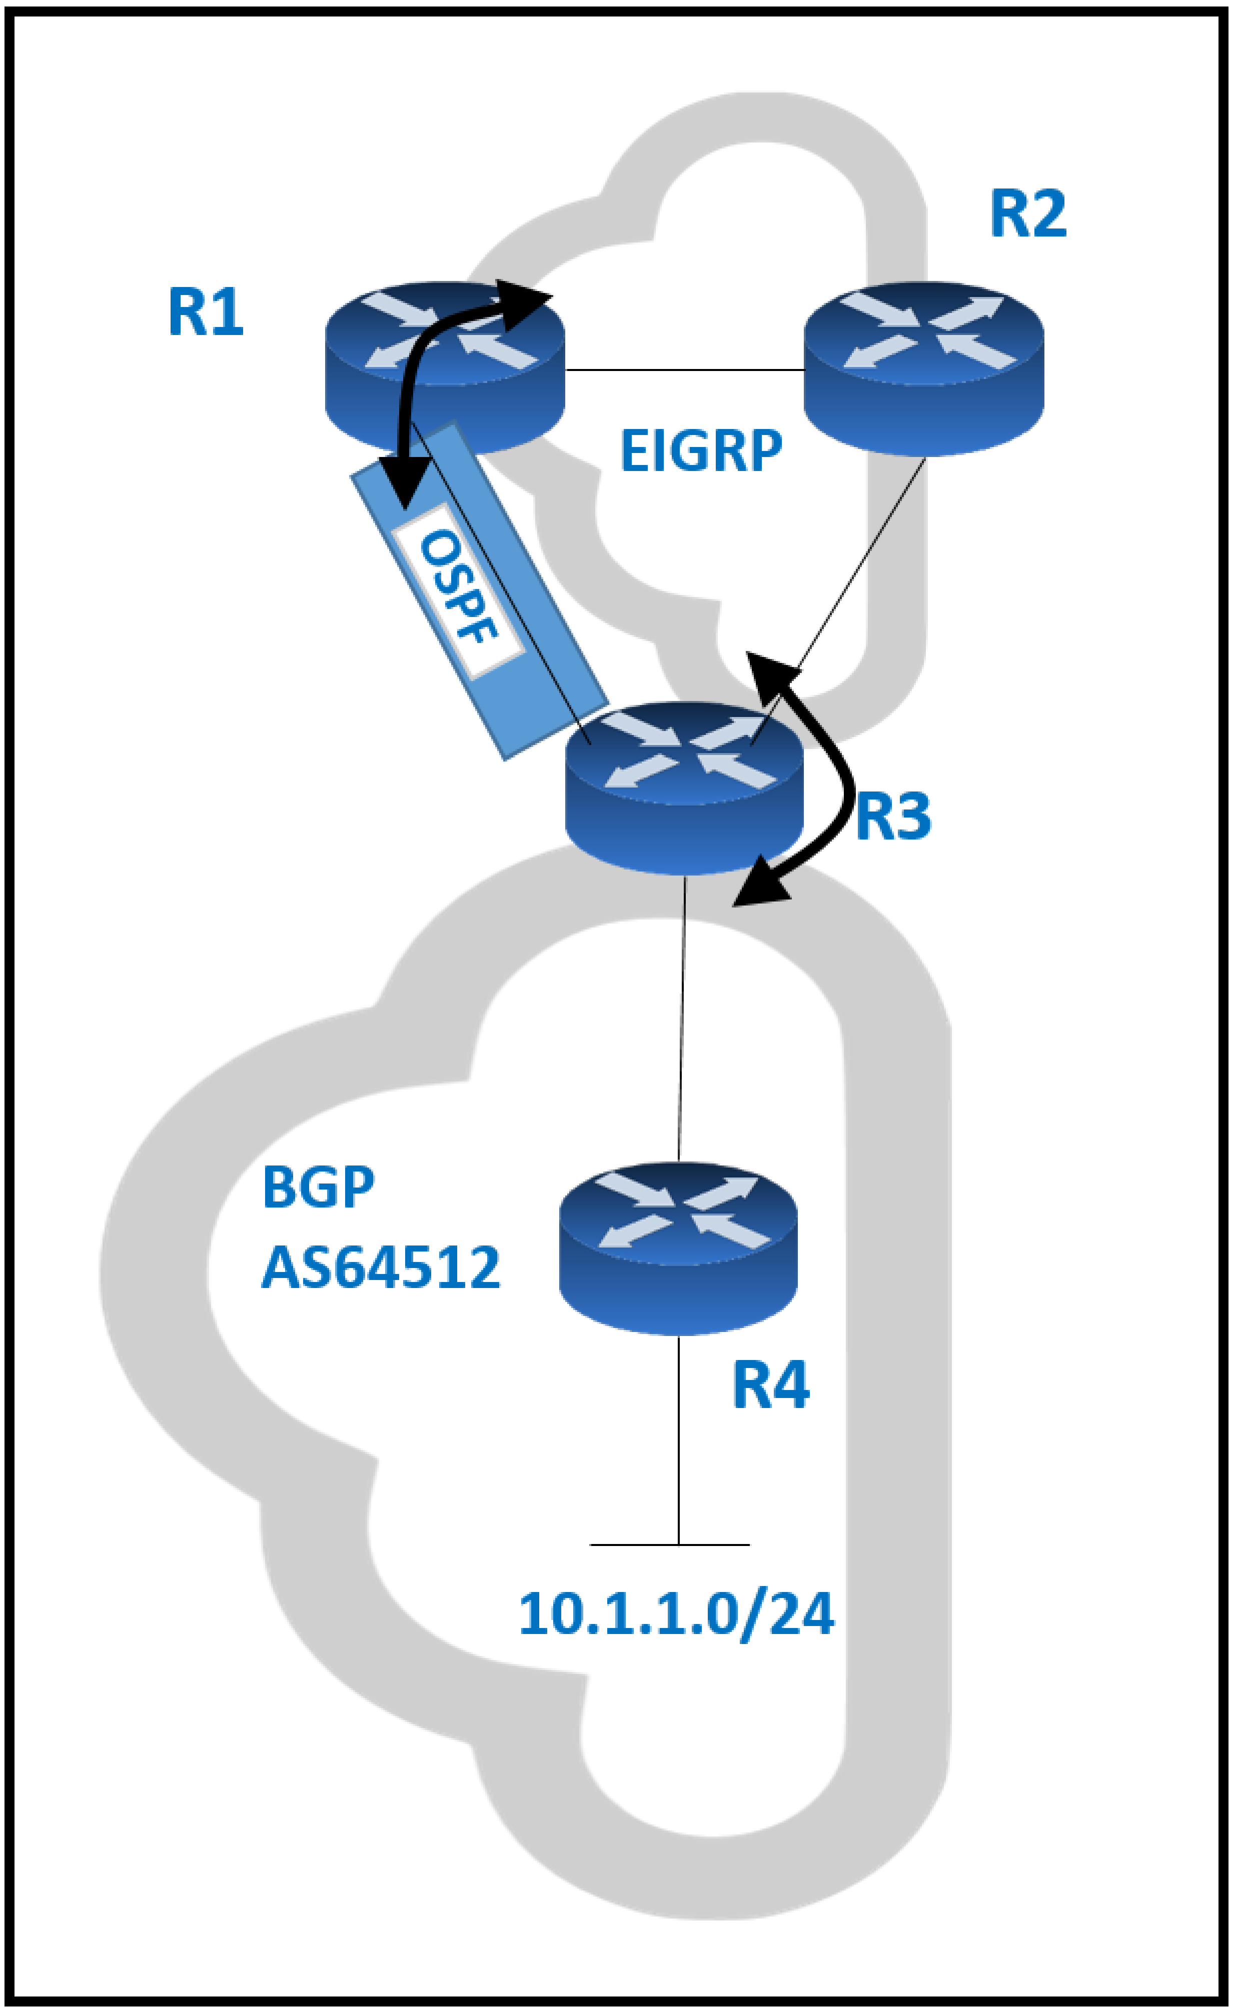

Refer to the exhibit. Routing protocols are mutually redistributed on R3 and R1. Users report intermittent connectivity to services hosted on the 10.1.1.0/24 prefix.

Significant routing update changes are noticed on R3 when the show ip route profile command is run.

How must the services be stabilized?

A. The routing loop must be fixed by reducing the admin distance of OSPF from 110 to 80 on R3

B. The routing loop must be fixed by reducing the admin distance of iBGP from 200 to 100 on R3

C. The issue with using BGP must be resolved by using another protocol and redistributing it into EIGRP on R3

D. The issue with using iBGP must be fixed by running eBGP between R3 and R4

Correct Answer: *B*

Community vote distribution

B (100%)

⊟ **AliMo123** Highly Voted 2 years, 1 month ago

B is correct
After redistribution, R3 learns about network 10.1.1.0/24 via two paths:+ Internal BGP (IBGP):
advertised from R4 with AD of 200 (and metric of 0)+ OSPF: advertised from R1 with AD of 110 (O E2) (and metric of 20)Therefore R3 will choose the path with the lower AD via OSPF But this is a looped path which is received from R3 -> R2 -> R1 -> R3. So when the advertised route from R4 is expired, the looped path is also expired soon and R3 willreinstall the main path from R4. This is the cause of intermittent connectivity.In order to solve this issue, we can lower the AD of iBGP to a value which is lower than 110 so that it is preferred over OSPF-advertised route.

upvoted 17 times

⊟ **wts** 1 year, 10 months ago

Routing protocols are mutually redistributed on R3 and R1. R3 learns about network 10.1.1.0/24 via three paths:
OSPF(110[O E2]),
EIGRP(170[D EX]),
iBGP(200).
The IBGP has too much administrative distance. Packets with a destination address from the 10.1.1.0/24 subnet miss this path and travel in a circle.

C and D are too strange options.
Reducing the administrative distance makes the routes of this protocol more preferable. 10.1.1.0/24 is behind the BGP, hence option B.

upvoted 1 times

⊟ **Jenia1** 1 year, 10 months ago

B should be correct,
C is general and we have more specific options.
D would be correct if they mentioned public AS/IP, but u can see private on the exhibit

upvoted 1 times

⊟ **[Removed]** Most Recent 4 months ago

Selected Answer: B

This is one lazy looking diagram.

upvoted 1 times

⊟ **Malasxd** 7 months, 2 weeks ago

Selected Answer: B

B is correct

upvoted 2 times

Question #42 *Topic 1*

When determining if a system is capable of support, what is the minimum time spacing required for a BFD control packet to receive once a control packet is arrived?

A. Desired Min TX Interval

B. Detect Mult

C. Required Min RX Interval

D. Required Min Echo RX Interval

Correct Answer: *C*

Community vote distribution

C (100%)

Networkingguy 1 year, 10 months ago

Selected Answer: C

Nice one ciscomicha, C is correct

upvoted 1 times

ciscomicha 1 year, 11 months ago

Selected Answer: C

C. Source: https://www.cisco.com/en/US/technologies/tk648/tk365/tk480/technologies_white_paper0900aecd80244005.html

upvoted 4 times

Stivostine 2 years ago

C is ok

Required Min RX Interval : This is the minimum interval, in microseconds, between received BFD Control packets that this system is capable of supporting.

upvoted 2 times

Question #43 *Topic 1*

An engineer is configuring a network and needs packets to be forwarded to an interface for any destination address that is not in the routing table.
What should be configured to accomplish this task?

A. set ip next-hop

B. set ip default next-hop

C. set ip next-hop recursive

D. set ip next-hop verify-availability

Correct Answer: *B*

Community vote distribution

B (100%)

Malasxd 7 months, 2 weeks ago

Selected Answer: B

B for sure

upvoted 1 times

ellen_AA 11 months, 3 weeks ago

- The "set ip default next-hop" command verifies the existence of the destination IP address in the routing table:
* If the destination IP address exists in the RT, the command does not policy route the packet, but forwards the packet based on the routing table.
* If the destination IP address does not exist in the RT, the command policy routes the packet by sending it to the specified next hop.

- The "set ip next-hop" command verifies the existence of the destination IP address in the routing table:
* If the next hop exists in the routing table, then the command policy routes the packet to the next hop.
* If the next hop does not exist in the routing table, the command uses the routing table to forward the packet.

upvoted 3 times

mrnipsnips 1 year, 1 month ago

Selected Answer: B

B for sure

upvoted 2 times

Stupid 1 year, 11 months ago

Selected Answer: B

Reference: https://www.cisco.com/c/en/us/support/docs/ip/ip-routed-protocols/47121-pbr-cmds-ce.html#:~:text=The%20set%20ip%20default%20next%2Dhop%20command%20verifies,by%20sending%20it%20to%20the%20specified%20next%20hop

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

Question #44 Topic 1

What is an advantage of using BFD?

A. It detects local link failure at layer 1 and updates the routing table.

B. It detects local link failure at layer 3 and updates the routing protocols.

C. It has sub-second failure detection for layer 1 and layer 3 problems.

D. It has sub-second failure detection for layer 1 and layer 2 problems.

Correct Answer: *D*

Community vote distribution

D (53%) B (47%)

Koume Highly Voted 11 months, 1 week ago

Selected Answer: D

I Go for D just following reasion
A. It detects local link failure at layer 1 and updates the routing table. WRONG
- BFD detects local link failures, but BFD does not interact with the routing table WRONG

B. It detects local link failure at layer 3 and updates the routing protocols. (ALMOST RIGHT BUT WRONG)
- The question here is What is the advantage of using BFD. routing procols can detect local link failures, so this is not an advantage.

C. It has sub-second failure detection for layer 1 and layer 3 problems. WRONG TRICKY'
- sub second failure detection is an advange of BFD, but BFD only detects L1-L2 problems, ink itself, BFD can not detect L3 problem like address misconfig.

D. It has sub-second failure detection for layer 1 and layer 2 problems. RIGHT!!
- The advantage of BFD is it's sub-second failure detection, and just detect L1-L2 problems.

upvoted 5 times

louisvuitton12 Most Recent 1 month, 3 weeks ago

Selected Answer: B

From Cisco U ENARSI Course:

"Typically, BFD can be used at any protocol layer. However, the Cisco implementation of BFD supports only Layer 3 clients, in particular, the BGP, Enhanced Interior Gateway Routing Protocol (EIGRP), IS-IS, and OSPF routing protocols, as well as the high availability protocol HSRP and also static routing."

upvoted 1 times

BTK0311 3 months ago

The advantage of using BFD (Bidirectional Forwarding Detection) is described as:

C. It has sub-second failure detection for layer 1 and layer 3 problems.

BFD provides rapid detection of network failures at both Layer 1 (physical layer) and Layer 3 (network layer), and it can detect these problems within milliseconds (sub-second). This quick detection helps in minimizing network downtime and improving network reliability by promptly identifying and responding to issues at these layers. per ChatGPT

upvoted 1 times

Fenix7 3 months, 1 week ago

It's B. Look at the text below from Cisco.

"BFD treats routing protocols, such as OSPF, as clients for creating the BFD sessions. The routing protocol discovers the neighbor using its own detection mechanism and then uses this information to form the BFD session with the neighboring router. If a link failure is detected by BFD, the client routing protocol is notified. This allows OSPF to tear down the routing neighbor adjacency immediately, instead of waiting multiple seconds for the hold timers to expire."

upvoted 1 times

Youssefmetry 3 months, 3 weeks ago

BFD can be used at any protocol layer. It could, for example, detect Physical or Data Link layers failures.
https://www.cisco.com/en/US/technologies/tk648/tk365/tk207/technologies_white_paper0900aecd80243fe7.html

upvoted 2 times

JieW 4 months, 2 weeks ago

Selected Answer: B

Voting B.
BFD is designed for IP level failure detection. Read section 2. Design. It does make notice of physical links but not the layer itself. It only refers to layers 2 and 3 by name.
https://datatracker.ietf.org/doc/html/rfc5880

upvoted 1 times

mabus 5 months, 1 week ago

Selected Answer: B

B is corect

upvoted 1 times

mabus 5 months, 3 weeks ago

Selected Answer: B

BFD works only for directly connected neighbors. BFD neighbors must be no more than one IP hop away. Multihop configurations are not supported -> "BFD detects local link failure" is correct.

Typically, BFD can be used at any protocol layer. However, the Cisco implementation of BFD for Cisco IOS Releases 12.2(18)SXE, 12.0(31)S, and 12.4(4)T supports only Layer 3 clients, in particular, the BGP, EIGRP, IS-IS, and OSPF routing protocols.

Reference: https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fs_bfd.html

According to the reference above, it is a bit weird but answer B is the best choice here.

upvoted 2 times

adudeguy 6 months, 2 weeks ago

B. BFD can be used at any protocol layer. It could, for example, detect Physical or Data Link layers failures, if the existing mechanisms did not provide sufficiently speedy detection. However, in the first phase of Cisco BFD support, all BFD clients, particularly the Layer 3 routing protocols (OSPF, IS-IS, EIGRP, and BGP) are at the Network layer.

upvoted 1 times

MasterMatt 8 months, 2 weeks ago

Selected Answer: B

Typically, BFD can be used at any protocol layer. However, the Cisco implementation of BFD for Cisco IOS Releases 12.2(18)SXE, 12.0(31)S, and 12.4(4)T supports only Layer 3 clients, in particular, the BGP, EIGRP, IS-IS, and OSPF routing protocols.

Source: https://www.cisco.com/en/US/docs/ios/12_4t/ip_route/configuration/guide/t_bfd.html

upvoted 1 times

6dd4aa0 8 months, 3 weeks ago

Selected Answer: B

Typically, BFD can be used at any protocol layer. However, the Cisco implementation of BFD for Cisco IOS Releases 12.2(18)SXE, 12.0(31)S, and 12.4(4)T supports only Layer 3 clients, in particular, the BGP, EIGRP, IS-IS, and OSPF routing protocols.

upvoted 1 times

DaanB 1 year, 1 month ago

Conclusion here https://www.cisco.com/en/US/technologies/tk648/tk365/tk207/technologies_white_paper0900aecd80243fe7.html says:
Bidirectional Forwarding Detection provides a method for network administrators to configure sub-second Layer 2 failure detection between adjacent network nodes. Furthermore, they can configure their routing protocols to respond to BFD notifications, and begin Layer 3 route convergence almost immediately.
If I understand this correctly, you (network administrator) use the method provided by BFD to configure Layer 2 sub-second failure detection. Then you can configure your routing protocol to act upon BFD notifications and begin route recalculations.

upvoted 1 times

Edwinmolinab 1 year, 5 months ago

Selected Answer: D

From https://www.cisco.com/en/US/technologies/tk648/tk365/tk207/technologies_white_paper0900aecd80243fe7.html
BFD can be used at any protocol layer. It could, for example, detect Physical or Data Link layers failures, if the existing mechanisms did not provide sufficiently speedy detection. However, in the first phase of Cisco BFD support, all BFD clients, particularly the Layer 3 routing protocols (OSPF, IS-IS, EIGRP, and BGP) are at the Network layer.

upvoted 1 times

balari 1 year, 5 months ago

B
https://www.cisco.com/en/US/docs/ios/12_4t/ip_route/configuration/guide/t_bfd.html

upvoted 1 times

IceFireSoul 1 year, 5 months ago

Correct answer B:
See https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fs_bfd.html

BFD provides fast BFD peer failure detection times independently of all media types, encapsulations, topologies, and routing protocols BGP, EIGRP, IS-IS, and OSPF. By sending rapid failure detection notices to the routing protocols in the local router to initiate the routing table recalculation process, BFD contributes to greatly reduced overall network convergence time. Figure 1 shows a simple network with two routers running OSPF and BFD. When OSPF discovers a neighbor (1) it sends a request to the local BFD process to initiate a BFD neighbor session with the OSPF neighbor router (2). The BFD neighbor session with the OSPF neighbor router is established (3).

upvoted 1 times

Dacusai 1 year, 6 months ago

B is the correct answer. The book for this exam talk only about detect routing problems nothing about layer 1 or 2, if a layer 1 issue occur the int goes down, and so with layer 2 protocol down.

upvoted 1 times

xziomal9 1 year, 7 months ago

Selected Answer: B

The correct answer is: B

upvoted 2 times

Question #45 *Topic 1*

An engineer needs dynamic routing between two routers and is unable to establish OSPF adjacency. The output of the show ip ospf neighbor command shows that the neighbor state is EXSTART/EXCHANGE.
Which action should be taken to resolve this issue?

A. match the passwords

B. match the hello timers

C. match the MTUs

D. match the network types

Correct Answer: *C*
Reference:
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13684-12.html

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #46 Topic 1

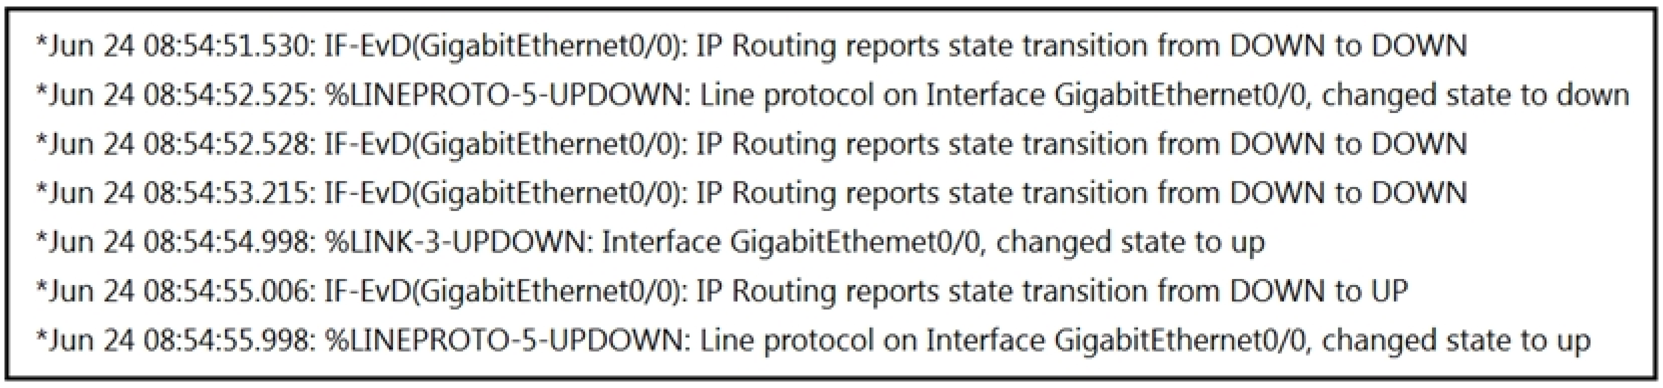

```
*Jun 24 08:54:51.530: IF-EvD(GigabitEthernet0/0): IP Routing reports state transition from DOWN to DOWN
*Jun 24 08:54:52.525: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/0, changed state to down
*Jun 24 08:54:52.528: IF-EvD(GigabitEthernet0/0): IP Routing reports state transition from DOWN to DOWN
*Jun 24 08:54:53.215: IF-EvD(GigabitEthernet0/0): IP Routing reports state transition from DOWN to DOWN
*Jun 24 08:54:54.998: %LINK-3-UPDOWN: Interface GigabitEthemet0/0, changed state to up
*Jun 24 08:54:55.006: IF-EvD(GigabitEthernet0/0): IP Routing reports state transition from DOWN to UP
*Jun 24 08:54:55.998: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/0, changed state to up
```

Refer to the exhibit. R1 is connected with R2 via GigabitEthernet0/0, and R2 cannot ping R1.

What action will fix the issue?

A. Fix route dampening configured on the router.

B. Replace the SFP module because it is not supported.

C. Fix IP Event Dampening configured on the interface.

D. Correct the IP SLA probe that failed.

Correct Answer: *C*

Mjestic Highly Voted 2 years, 3 months ago

For those like me who don't what is IP Event Dampening :

The IP Event Dampening feature introduces a configurable exponential decay mechanism to suppress the effects of excessive interface flapping events on routing protocols and routing tables in the network. This feature allows the network operator to configure a router to automatically identify and selectively dampen a local interface that is flapping. Dampening an interface removes the interface from the network until the interface stops flapping and becomes stable. Configuring the IP Event Dampening feature improves convergence times and stability throughout the network by isolating failures so that disturbances are not propagated. This, in turn, reduces the utilization of system processing resources by other devices in the network and improves overall network stability.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/xe-16-11/iri-xe-16-11-book/iri-pi-event-damp.html

upvoted 26 times

Eddyyin 5 months, 3 weeks ago

Would you please explain more, how can this feature fix the issue? Isn't the real issue that needs fixing is the flapping port itself?

upvoted 1 times

Almylle Highly Voted 6 months, 2 weeks ago

Why cisco create this type of questions, i read completely the Enarsi book from cisco press and this is not writed in that book.

upvoted 8 times

ledesir 2 weeks, 5 days ago

sale for me , never heard about it before , its not even in the enarsi book

upvoted 1 times

error_909 Most Recent 2 years, 3 months ago

The given answer is correct

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 3 times

Question #47 *Topic 1*

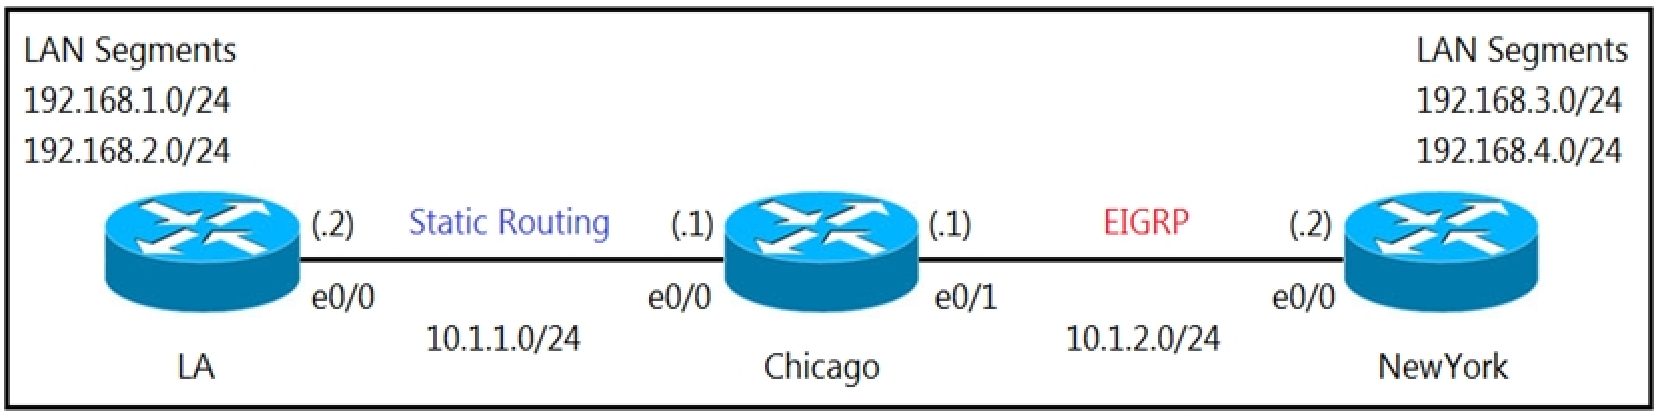

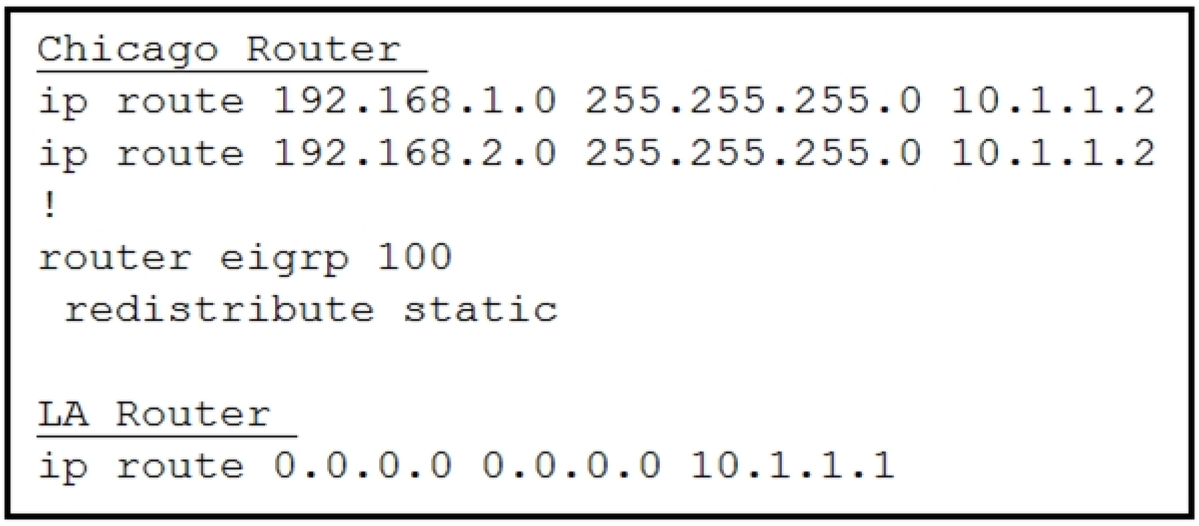

```
Chicago Router
ip route 192.168.1.0 255.255.255.0 10.1.1.2
ip route 192.168.2.0 255.255.255.0 10.1.1.2
!
router eigrp 100
 redistribute static

LA Router
ip route 0.0.0.0 0.0.0.0 10.1.1.1
```

Refer to the exhibits. A user on the 192.168.1.0/24 network can successfully ping 192.168.3.1, but the administrator cannot ping 192.168.3.1 from the LA router.

Which set of configurations fixes the issue?

A.

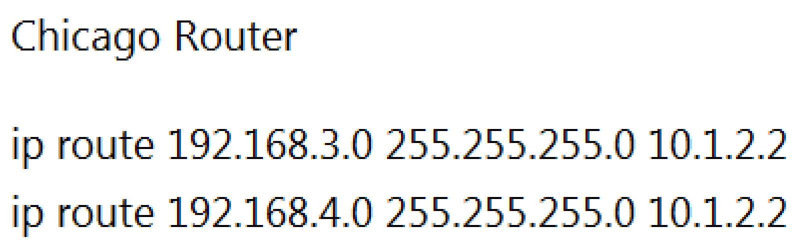

```
Chicago Router

ip route 192.168.3.0 255.255.255.0 10.1.2.2
ip route 192.168.4.0 255.255.255.0 10.1.2.2
```

B.

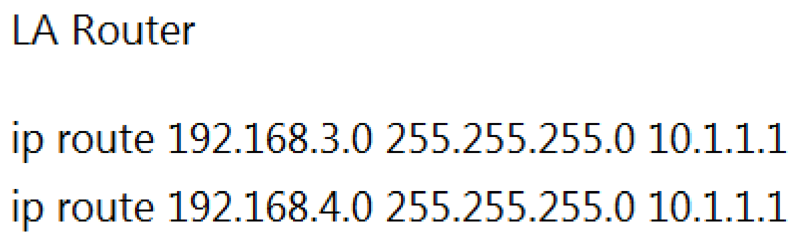

```
LA Router

ip route 192.168.3.0 255.255.255.0 10.1.1.1
ip route 192.168.4.0 255.255.255.0 10.1.1.1
```

C.

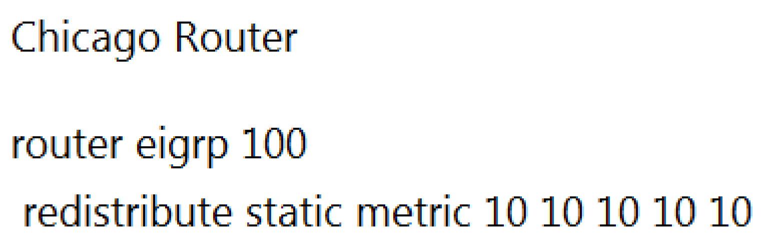

```
Chicago Router

router eigrp 100
 redistribute static metric 10 10 10 10 10
```

D.

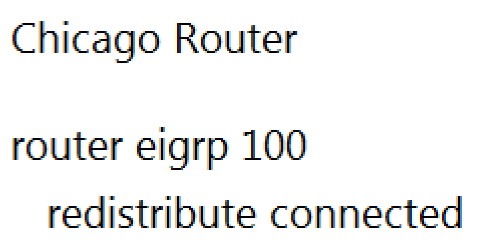

```
Chicago Router

router eigrp 100
  redistribute connected
```

Correct Answer: *D*

⊟ **xqlz** Highly Voted 👍 2 years, 1 month ago

Correct answer is D and not C

The administrator is isuing the ping from LA router, so the source IP will be 10.1.1.2
So when the reply comes back from 192.168.3.1 the destination will be 10.1.1.2 but the NewYork router doesn't have a route for that destination.

If the redistribute static (without the metric) was not working then the first ping would also fail since NewYork router would not have a route to 192.168.1.0/24

Please correct if wrong.
upvoted 14 times

HungarianDish Most Recent 7 months ago
For me, it is simply "D". Based on the topology, the networks 192.168.3.0/24 and 192.168.4.0/24 belong to the eigrp domain. Thus, "redistribute connected" under eigrp process is enough to provide connectivity from LA to NY. (I also confirmed it in a lab.)
upvoted 4 times

MasterMatt 8 months, 2 weeks ago
Emulated this in the lab and while eigrp has a redistribute static will advertise the external route to New York router, once the ICMP is sent back to New York from Chicago it is dropped as we don't have a route entry for the 192.168.3.* and 192.168.4.*. It says that with the initial config it works but it don't. There is no point in redistributing connected and no need of adding the static matrics. For me the correct answer is A, if we need the pings to work.
upvoted 3 times

> **David98898998** 6 months, 3 weeks ago
> Because the ping from 192.168.1.0 works, this implies that NY has static routes routes to it and is sharing them with Chicago.
>
> The fact that the admins ping from LA doesn't work, implies that NY isn't aware of the 10.1.1.0 network. NY doesn't need to be aware of this network to reach the LAN networks off of LA, but to reach LA router itself, it must be made aware. This can be done by redistributing connected routes on Chicago.
> upvoted 2 times

6dd4aa0 8 months, 3 weeks ago
The question does not specify if EIGRP is configured for R3 on the network for 192.168.3.0 and 192.168.4.0

Assuming EIGRP is configured for 192.168.3.0 and 192.168.4.0
==
The correct Answer is D.

Why not A? Because in R3, it is configured with
EIGRP 100
network 192.168.2.0 0.0.0.255
network 192.168.3.0 0.0.0.255

As a result, EIGRP will populate these 2 routes to R2. Hence, configuring a static route will do the trick, it defeats the purpose of EIGRP. Moreover, the static routes will have an AD of 1, which will then overwrite Eigrp AD of 90.

Assuming EIGRP is NOT configured for 192.168.3.0 and 192.168.4.0
==
The correct Answer is A.
upvoted 1 times

forccnp 9 months, 4 weeks ago
D is the correct answer
upvoted 1 times

Router 1 year, 3 months ago
c is the correct ans, you must specify metric if you're redistributing into eigrp
upvoted 1 times

johnmcclane78 1 year, 5 months ago
Correct answer is A. Tested in lab.
B - wrong next-hop
C - doesn't make sense, static routes will be available without metric too
D - it changes nothing.

The problem is absense of routes on Chicago to 3.0/24 and 4.0/24. That's why ping doesn't work. And A is the only way to fix it (except "redistribute connected" into EIGRP on NY)
upvoted 1 times

Hack4 1 year, 10 months ago
After doing this lab, i think " Redistribute connected" is most appropriate....Because by doing 'redistributed connected routes", the 10.1.1.0/24 is being seen as EIGRP external route by the New-York router..
upvoted 1 times

kent2612 1 year, 10 months ago
Ans should be A & D
I lab it up, redistribute connected alone (on Chicago) don't work since Chicago didn't know how to reach 192.168.3.0/24 and 192.168.4.0/24
upvoted 1 times

> **AliMo123** 1 year, 10 months ago
> D is correct
> see the ip route 0.0.0.0 0.0.0.0 10.1.1.1 on LA router which enables Chicago router to route from LA to NY routers
> upvoted 1 times

Carl1999 1 year, 10 months ago

D is correct.
New york dosent know 10.1.1.0/24.

C commands are for OSPF not EIGRP.

upvoted 2 times

geek1992 1 year, 11 months ago

Why not A ?

upvoted 1 times

[Removed] 1 year, 11 months ago

Well since the user can ping the 192.168.3.0 network, why would you place a static route to reach those networks? The issue is the traffic coming back. The admin ping is making it to the network but the traffic isnt coming back because there isnt a route back to the admin. Thats why you redistribute connected into EIGRP.

upvoted 2 times

geek1992 1 year, 11 months ago

C is correct seed metric eigrp is infinity

upvoted 2 times

Carl1999 1 year, 10 months ago

A seed metric of 1 is given when redistributed from connected and static routing processes.
So, if the redelivery source is connected or static, you do not need to set the seed metric.

upvoted 1 times

JOKERR 2 years ago

I think D is correct.

For redistribute connected, EIGRP will have a look at the component metrics (bandwidth, delay, optionally reliability and load) of the interfaces where these networks are connected, and will compute the resulting metric out of these values. This is, by the way, precisely the same thing EIGRP would do if you had the networks added by a network command.

https://community.cisco.com/t5/switching/eigrp-redistribute-connected/td-p/2878396

upvoted 1 times

mosvan 2 years, 4 months ago

Chicago router is only missing the connected route. So the simplest solution would be answer D

upvoted 4 times

RS_nw 2 years, 4 months ago

C is the correct Answer.

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

Question #48 Topic 1

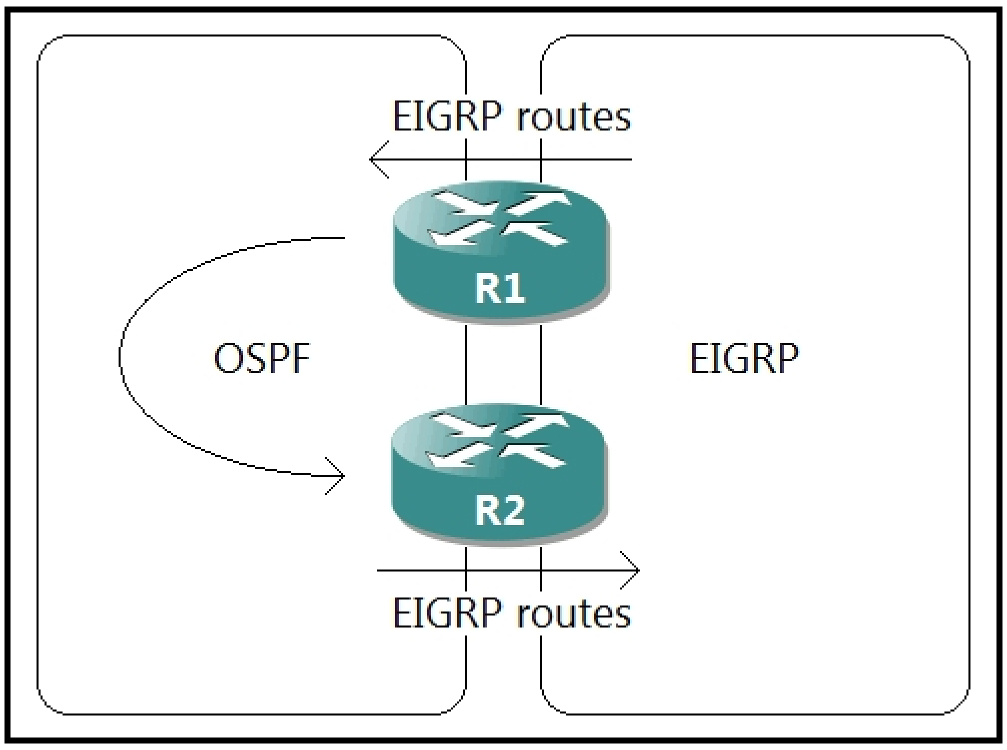

Refer to the exhibit. A network administrator configured mutual redistribution on R1 and R2 routers, which caused instability in the network. Which action resolves the issue?

A. Set a tag in the route map when redistributing EIGRP into OSPF on R1, and match the same tag on R2 to deny when redistributing OSPF into EIGRP.

B. Set a tag in the route map when redistributing EIGRP into OSPF on R1, and match the same tag on R2 to allow when redistributing OSPF into EIGRP.

C. Apply a prefix list of EIGRP network routes in OSPF domain on R1 to propagate back into the EIGRP routing domain.

D. Advertise summary routes of EIGRP to OSPF and deny specific EIGRP routes when redistributing into OSPF.

Correct Answer: *A*

Community vote distribution

A (100%)

Malasxd 7 months, 2 weeks ago

Selected Answer: A

the given answer is correct

upvoted 1 times

DumpsterFire 1 year, 3 months ago

Selected Answer: A

A is correct.

upvoted 1 times

Hack4 1 year, 10 months ago

the given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

Question #49 Topic 1

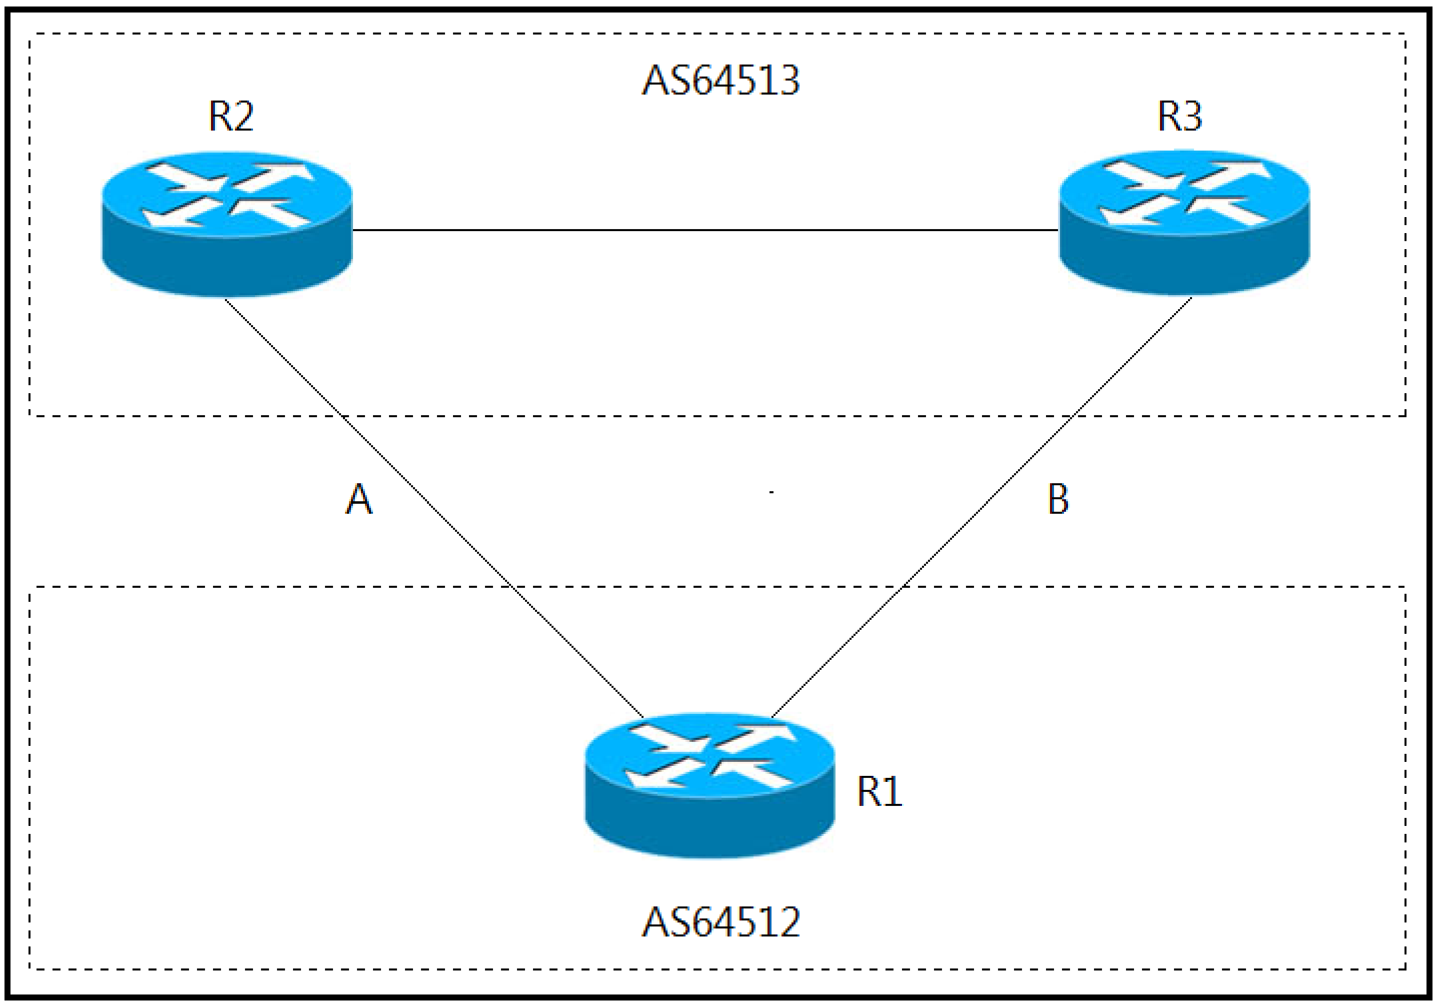

Refer to the exhibit. A network engineer for AS64512 must remove the inbound and outbound traffic from link A during maintenance without closing the BGP session so that there is still a backup link over link A toward the ASN.

Which BGP configuration on R1 accomplishes this goal?

A.

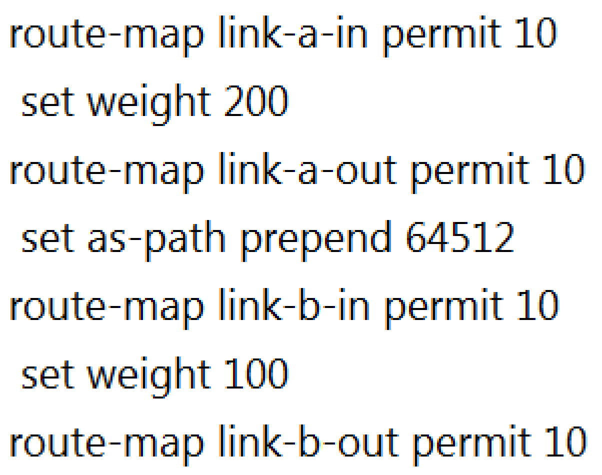

```
route-map link-a-in permit 10
 set weight 200
route-map link-a-out permit 10
 set as-path prepend 64512
route-map link-b-in permit 10
 set weight 100
route-map link-b-out permit 10
```

B.

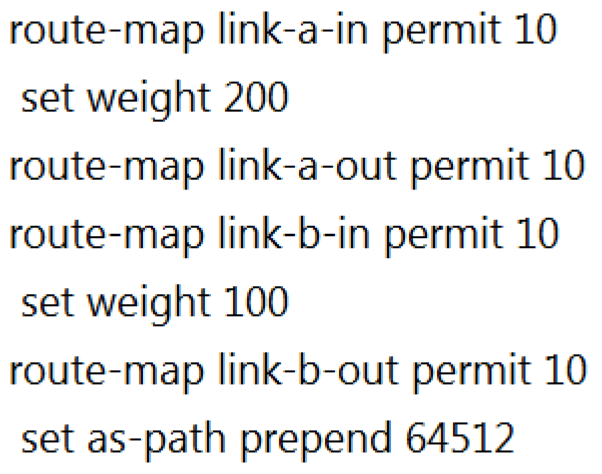

```
route-map link-a-in permit 10
 set weight 200
route-map link-a-out permit 10
route-map link-b-in permit 10
 set weight 100
route-map link-b-out permit 10
 set as-path prepend 64512
```

C.

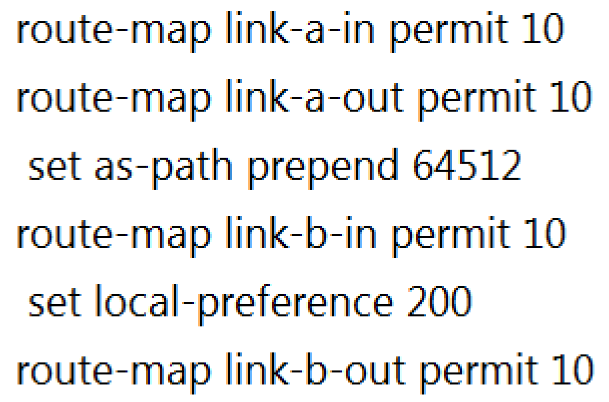

```
route-map link-a-in permit 10
route-map link-a-out permit 10
 set as-path prepend 64512
route-map link-b-in permit 10
 set local-preference 200
route-map link-b-out permit 10
```

D.

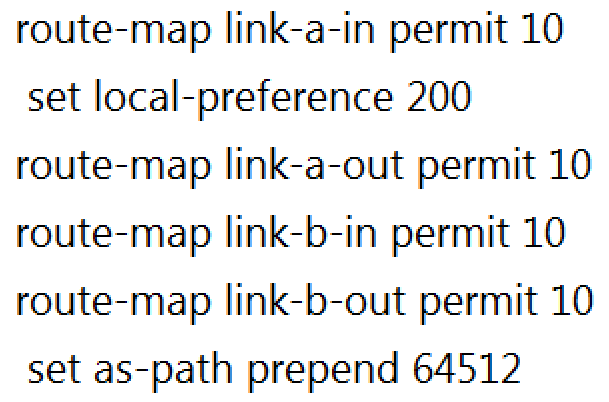

```
route-map link-a-in permit 10
 set local-preference 200
route-map link-a-out permit 10
route-map link-b-in permit 10
route-map link-b-out permit 10
 set as-path prepend 64512
```

Correct Answer: *C*

inteldarvid 5 months, 2 weeks ago

the option corret is C

upvoted 1 times

XBfoundX 5 months, 4 weeks ago

Hello,

the only one correct here is C because the other one are making the incoming routes from R2 to be better than the routes of router R3. In this case cause we have only one router we can put the weight 100 inbound for the updates coming from R3 (via a route-map that is going to set the weight to 100, by def is 0), then we configure AS prepend in R1 in outgoing direction for force R2 to go via R3 for reach R1 because the router see two organization instead of one to traverse.

upvoted 1 times

XBfoundX 5 months, 4 weeks ago

So in this case we are going to set the local preference inbound and outbound the as-prepend so answer is C!

upvoted 1 times

Wooker 9 months, 2 weeks ago

The given answer is correct "C"

upvoted 1 times

AliMo123 1 year, 9 months ago

none of them is correct
you do not need all these fancy route-map links to meet the solution
only route-map we need is route-map link-b to increase LP. that is all
create access-list
create route-map
match ip add
set the LP
apply the route-map to link b BGP

upvoted 2 times

wts 1 year, 9 months ago

LP will correct outgoing traffic. What will happen to the incoming?

upvoted 1 times

diogodds 1 year, 9 months ago

That is not really true, with that you will only be influencing the outbound traffic, what about the inbound?

upvoted 3 times

Hack4 1 year, 10 months ago

Yes correct

upvoted 1 times

GReddy2323 2 years ago

Can someone kindly explain what this is doing? Is this making the traffic start to take Link B while traffic A undergoes maintenance? Also, what is the purpose of "set as-path prepend 64512" for link A? Isn't local preference enough? Please more input is appreciated.

upvoted 1 times

JOKERR 2 years ago

That is correct. Config makes Traffic take link B while A is under maintenance.

route-map link-a-in permit 10 --> Allow incoming routes from A
route-map link-a-out permit 10 --> Prepend own AS-PATH so that AS-PATH becomes longer (for R2) to influence inbound traffic.
route-map link-b-in permit 10
set local-preference 200 --> Set local pref for routes coming in from B. Default local pref is 100 so R1 will choose B over A.
route-map link-b-out permit 10 --> permit all routes advertised to B.

Hope this clears. Please comment any mistakes/corrections.

upvoted 9 times

Stupid 1 year, 11 months ago

It´s making R1 think that incoming and outcoming traffic from link A is taking an extra hop, therefor making the as-path longer for R2, reference to this links "We will now use a technique called AS path prepending, which consists of adding extra “fake” hops to a path using our ASN multiple times."
https://nsrc.org/workshops/2018/ubuntunet-nren-bgp/networking/nren/en/labs/bgp-policy-as-prepend.html#:~:text=We%20will%20now%20use%20a%20technique%20called%20AS%20path%20prepending%2C%20which%20consists%20of%20adding%20extra%20%E2%80%9Cfake%E2%80%9D%20hops%20to%20a%20path%20using%20our%20ASN%20multiple%20times.

and https://www.ccexpert.us/routing-switching/step-4-shortest-aspath.html#:~:text=The%20concept%20and,not%20be%20intended.

You can check a configuration example here https://community.cisco.com/t5/networking-blogs/bgp-as-path-prepending-configuration/ba-p/3819334

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

THe given answer is correct
LP default is 100. Highest wins.

upvoted 1 times

Question #50 *Topic 1*

An engineer configured access list NON-CISCO in a policy to influence routes.

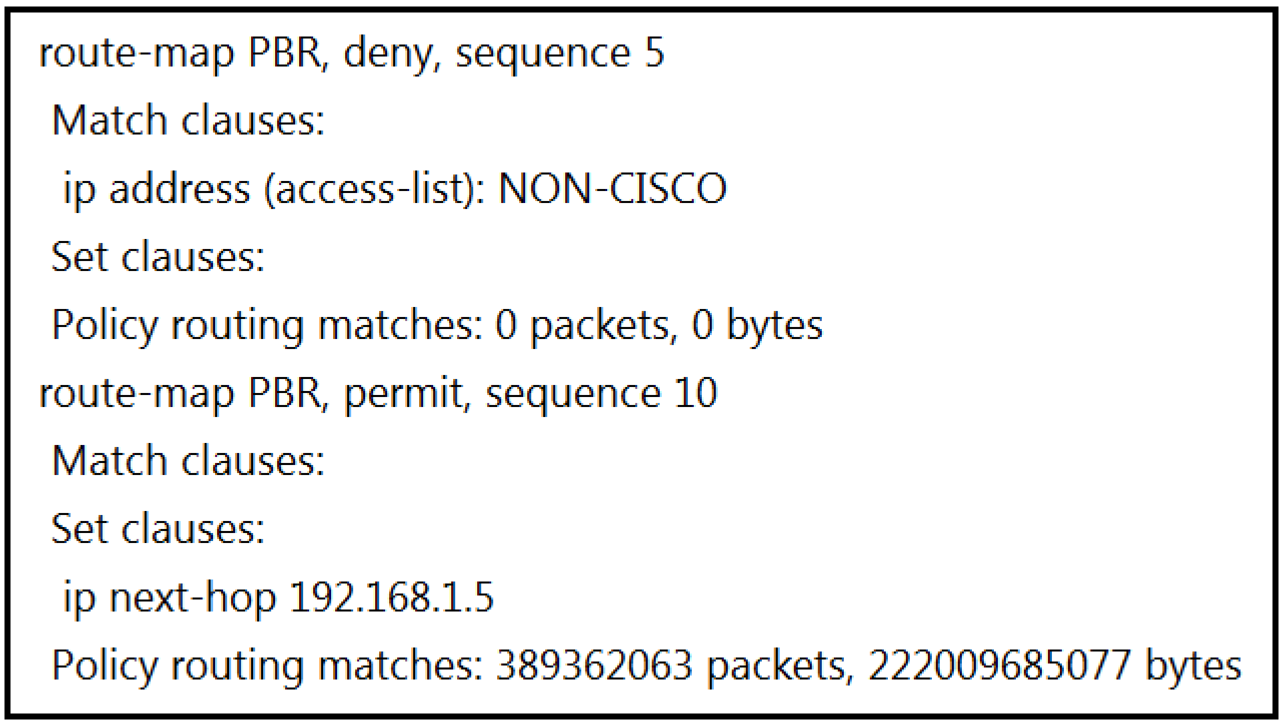

```
route-map PBR, deny, sequence 5
 Match clauses:
  ip address (access-list): NON-CISCO
 Set clauses:
 Policy routing matches: 0 packets, 0 bytes
route-map PBR, permit, sequence 10
 Match clauses:
 Set clauses:
  ip next-hop 192.168.1.5
 Policy routing matches: 389362063 packets, 222009685077 bytes
```

What are the two effects of this route map configuration? (Choose two.)

A. Packets are forwarded using normal route lookup.

B. Packets are forwarded to the default gateway.

C. Packets are dropped by the access list.

D. Packets are evaluated by sequence 10.

E. Packets are not evaluated by sequence 10.

Correct Answer: *BD*

Community vote distribution

AD (62%) AE (23%) CD (15%)

⊟ **ytsionis** Highly Voted 2 years ago

Seq 5 has a match ACL ---Deny
Seq 10 has no match so Match Everything ---Permit
So a packet
ether it matched by ACL and forwarded using normal route lookup
or does not get matched by ACL and evaluated by sequence 10.
A , D

upvoted 18 times

⊟ **Pietjeplukgeluk** 1 month, 1 week ago

I agree on A and D being "the best options" however, please note either A or B are both incorrect in a way, A states " Packets are forwarded using normal route lookup", i think think it is to vague as the solution will just forward all packets to 192.168.1.5 (if in routing table). Anyway, i think this question is not of high quality.

upvoted 1 times

⊟ **fortinet1234** 2 months, 2 weeks ago

Since sequence 10 has no match condition that means that we can not evaluate according sequence 10 - So I guess the best options here are A & E

upvoted 1 times

⊟ **JOKERR** 1 year, 6 months ago

Yes. Makes sense. Thank you.

upvoted 1 times

⊟ **WAKIDI** 1 year, 5 months ago

sorry for my poor english. seq 10 has no match. Can we say seq 10 do an "evaluate" ?

upvoted 1 times

⊟ **YaPet** Highly Voted 1 year, 10 months ago

In my opinion B,D are correct answers.
No any packets are evaluated by seq 5. It means that all packets are evaluated by seq 10. Because it has permit statement and no match any conditions all packets are routed to 192.168.1.5 by PBR.
According to Cisco PBR command set-ip next hop explanation
The set ip next-hop command verifies the existence of the next hop specified, and...

... if the next hop exists in the routing table, then the command policy routes the packet to the next hop.
... if the next hop does not exist in the routing table, the command uses the normal routing table to forward the packet.
As we can see from output packets have been forwarded by sequence 10 and this is NO normal routing table. But here we need to be sure that 192.168.1.5 is default-gateway and it exists in the routing table.

upvoted 11 times

asans Most Recent 4 days, 20 hours ago

A and D
Any routes that match the NON-CISCO acl will be "denied", i.e. not processed by PBR and so will use the Routing Table (normal route lookup). =======> A

Any routes that do NOT match the NON-CISCO acl are permitted by seq 10 and thus use the Next-hop of 192.168.1.5 ======> D

upvoted 1 times

Ll123123 2 months ago

Selected Answer: AE

A E - because the seq 5 deny route map statement already mean the phr shall skipped to use routing table, so seq 10 is not evaluated. Tricky part is that it has matches for pbr matching because matching seq 5 is a match

upvoted 1 times

Ll123123 2 months ago

I will go with ae… I think the first deny in routemap already mean use routing table route in pbr. Pbr only execute upon a permit route map statement and has an implicit deny at the end. Since deny seq is before the permit, I think permit 10 won't be executed.. but better verify with simulator

upvoted 1 times

chris110 3 months, 2 weeks ago

Selected Answer: AD

Its A, D

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: AD

AD is optioN correct

upvoted 1 times

guy276465281819372 5 months, 3 weeks ago

Selected Answer: AD

A & D are correct.
either the packets are forwarded normally if they match the ACL else they are evaluated by sequence 10.

upvoted 1 times

XBfoundX 5 months, 4 weeks ago

As ytsionis says because the route-map do not have an acl that is matching the traffic the PBR will not be applied to any prefix because without the ACL the PBR is not gonna math nothing

upvoted 1 times

Malasxd 7 months ago

"A" and "D" are right.
If the packet match in ACL NON-CISCO, the route-map sequence 5 is set to deny it, but it is a PBR and not a filter, so the deny says to the packet follow the normal RIP lookup.
Any other packet that does not match NON-CISCO ACL will match here, so it will forwarded to 192.168.1.5.

upvoted 3 times

Titini 10 months ago

A &D As Jokerr mentioned. As we see we have hits only on route map 10 sequence, so we have D from that and what does this PBR sequence do?
b If you do not match packets on a route-map during PBR (as sequence 10), PBR does not take any action on that packet, and is routed normally per the routing table/FIB/etc. So we have A from there. (https://learningnetwork.cisco.com/s/question/0D53i00000Kt0jACAR/policy-based-routing)

upvoted 1 times

Lilienen 10 months ago

Selected Answer: AD

A and D

upvoted 1 times

tseen 10 months, 3 weeks ago

Selected Answer: CD

C. Packets are dropped by the access list.
D. Packets are evaluated by sequence 10.

upvoted 2 times

TheBaja 1 year, 1 month ago

The question is for packets that match ACL. For that packet, packets are evaluated in seq 5, and using normal route lookup. So my answare is A (normal route lookup) and E (not matched by sequence 10).

upvoted 1 times

Router 1 year, 3 months ago

a and d, packet that are denied will not be drop but be process by normal routing table and packets that a matched will be evaluated and forwarded to the next-hop

upvoted 2 times

Huntkey 1 year, 3 months ago

Selected Answer: AE
I am assuming the question specifically asks for packet being matched by the ACL, rather than other packets

upvoted 1 times

Edwinmolinab 1 year, 5 months ago

Selected Answer: AE

I don't see any matching clauses then to me the traffic would be use the normal route lookup and the packets won't be evaluated by sequence 10.

upvoted 1 times

Question #51 Topic 1

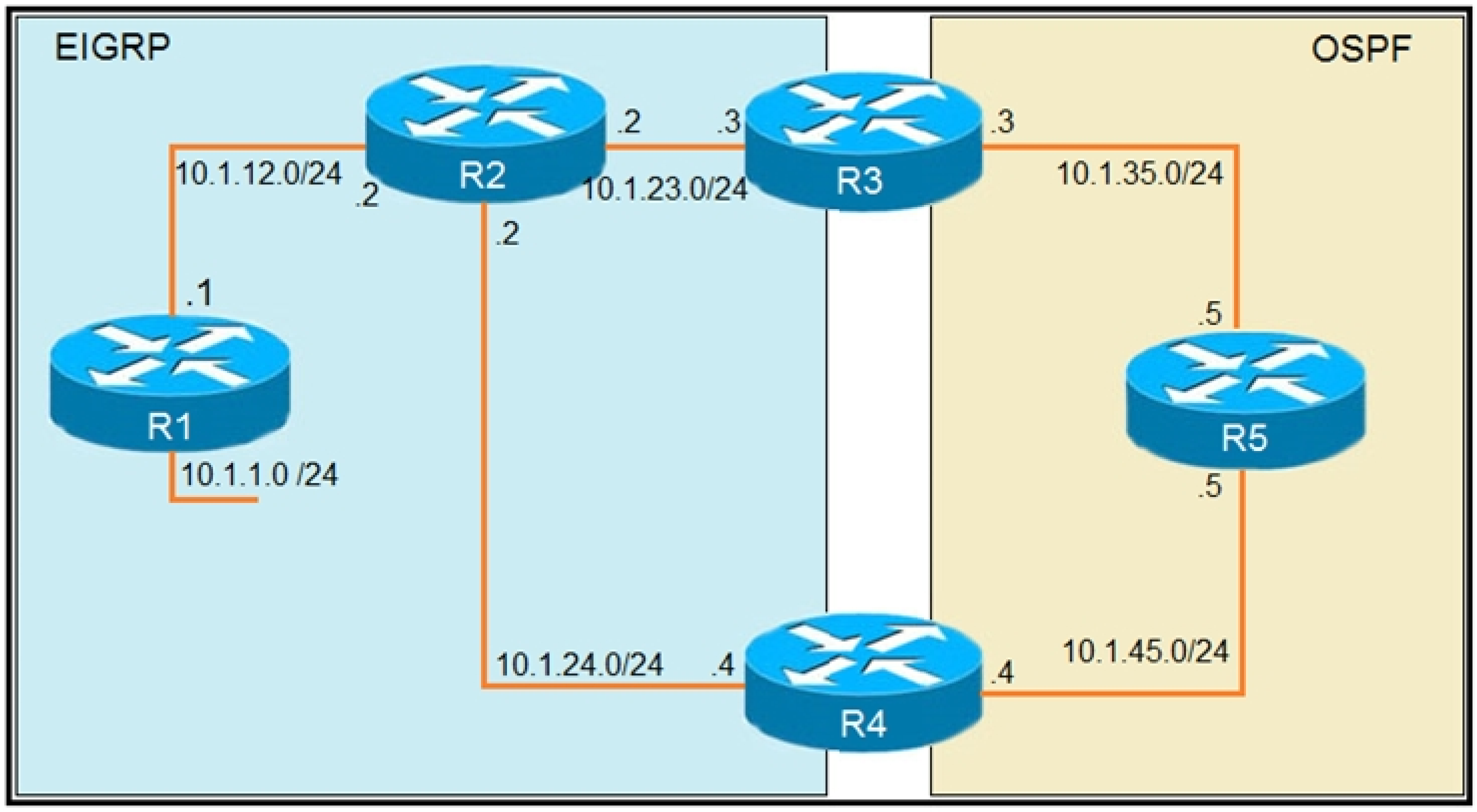

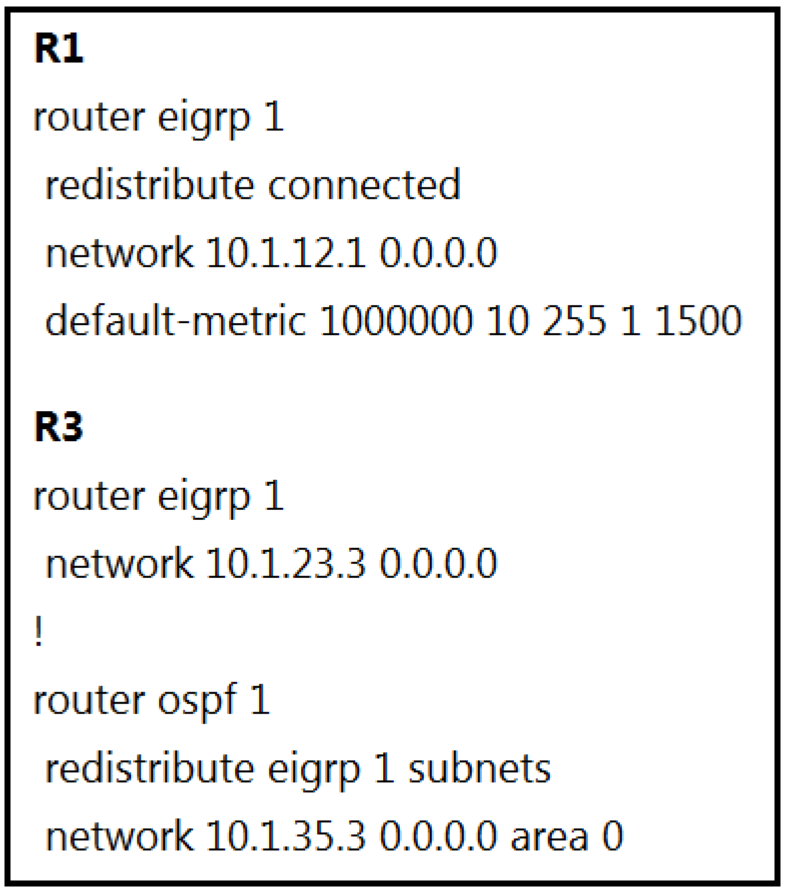

```
R1
router eigrp 1
 redistribute connected
 network 10.1.12.1 0.0.0.0
 default-metric 1000000 10 255 1 1500

R3
router eigrp 1
 network 10.1.23.3 0.0.0.0
!
router ospf 1
 redistribute eigrp 1 subnets
 network 10.1.35.3 0.0.0.0 area 0
```

Refer to the exhibits. To provide reachability to network 10.1.1.0/24 from R5, the network administrator redistributes EIGRP into OSPF on R3 but notices that R4 is now taking a suboptimal path through R5 to reach 10.1.1.0/24 network.
Which action fixes the issue while keeping the reachability from R5 to 10.1.1.0/24 network?

A. Change the administrative distance of the external EIGRP to 90.

B. Apply the outbound distribution list on R5 toward R4 in OSPF.

C. Change the administrative distance of OSPF to 200 on R5.

D. Redistribute OSPF into EIGRP on R4.

Correct Answer: *A*

Community vote distribution

A (100%)

Hammad745 Highly Voted 2 years, 4 months ago

The subnet 10.1.1.1/24 is redistributed into EIGRP domain so it will have the Administrative Distance (AD) of 170. Therefore R4 also learns about this subnet advertised from R2 with the same AD of 170.In the other hand, subnet 10.1.1.0/24 is also redistributed into OSPF on R3 so R5 & R4 will learn about this subnet with AD of 110, which is better than the above AD of 170 so R4 will choose path R4 -> R5 -> R3 -> R2 -> R1.
In order to solve this problem, we can configure an outbound distribute list on R5 to prevent (filter out) this subnet from advertising to R4. Then R4 only has one way to reach R1, which is R4 -> R2 -> R1. But this method will remove the backup route so it is not the best solution.Another solution is to reduce the AD of the external EIGRP to a value smaller than 110. This method reserves the backup route in case of the main route fails

upvoted 17 times

[Removed] 1 year, 11 months ago

B isnt correct.. You can't filter LSA's within the area. If it was inbound then yes but outbound are applied to the ABR/ASBR.

upvoted 3 times

JingleJangus 1 year, 11 months ago

All of the routers in the same OSPF area need to have the same exact LSDB. You cannot have it any other way in ospf. So this answer is wrong. You could implement a local distribute list on R4 to filter it locally from the RIB (still being in the LSDB), but I dont think this is THE BEST fix. The best fix is to maintain the backup route thru R5 and just lower the AD of external EIGRP to anything below 110.

upvoted 1 times

bryaberson Most Recent 4 months, 3 weeks ago

Why not C?

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: A

```
R4#trac 10.1.1.1
Type escape sequence to abort.
Tracing the route to 10.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
1 10.1.24.2 2 msec 2 msec 2 msec
2 10.1.12.1 2 msec * 2 msec
R4#sh run | sec router eigrp
router eigrp 1
network 10.1.24.0 0.0.0.255
redistribute ospf 1 metric 1000000 1 255 1 1500 route-map FILTER-TAG
distance eigrp 90 90
R4#
```

upvoted 1 times

wts 1 year, 10 months ago

Selected Answer: A

There are two routes on P4:
D EX (ad170) towards R2
O E2 (ad110) towards R5 (suboptimal path)

What needs to be done so that the packet goes towards R2, if a smaller administrative distance is preferable.

upvoted 1 times

wts 1 year, 9 months ago

It's external because it got into the EIGRP-domain through redistribute connected command.

upvoted 2 times

Carl1999 1 year, 10 months ago

Selected Answer: A

A is correct,

upvoted 1 times

JingleJangus 1 year, 11 months ago

Selected Answer: A

All of the routers in the same OSPF area need to have the same exact LSDB. You cannot have it any other way in ospf. So B is wrong. You could implement a local distribute list on R4 to filter it locally from the RIB (still being in the LSDB), but I dont think this is THE BEST fix. The best fix is to maintain the backup route thru R5 and just lower the AD of external EIGRP to anything below 110.

upvoted 1 times

LaughingGor 2 years, 2 months ago

B is right， after “”redistributes EIGRP into OSPF on R3 “,the administrative distance of "10.1.1.0/24 network"in ospf is 110. R4 has 2 path2 to "10.1.1.0/24 network":
R4-->R2: eigrp AD 170(because it is from "redistribute connected"commad on R1)
R4-->R5: OSPF AD 110(all AD value is "110 "in ospf including redistributed)
So it will choose R5-R4
B has no effect for LSA,right?
A is right:

```
R4(config)#router eigrp 1
R4(config-router)#distance eigrp 90 90
```

upvoted 3 times

error_909 2 years, 2 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct
external eigrp 170 ------> 90
ospf 110 ------> 110
the route is external eigrp because it is redistributed in.

upvoted 2 times

Precission21 2 years, 7 months ago

A is correct, 10. subnet is redistributed to eigrp so its treated as external route with defaul AD of 170

upvoted 2 times

RHK0783 2 years, 7 months ago

B is the correct answer. Make sure that R4 is already having the internal EIGRP route through R2. Outbound distribute-list on R5 is required to stop exporting the external routes to R4 i.e. learned from R3.

upvoted 2 times

[Removed] 1 year, 11 months ago

B isnt correct.. You can't filter LSA's within the area. If it was inbound then yes but outbound are applied to the ABR/ASBR.

upvoted 1 times

JingleJangus 1 year, 11 months ago

All of the routers in the same OSPF area need to have the same exact LSDB. You cannot have it any other way in ospf. So this answer is wrong. You could implement a local distribute list on R4 to filter it locally from the RIB (still being in the LSDB), but I dont think this is THE BEST fix. The best fix is to maintain the backup route thru R5 and just lower the AD of external EIGRP to anything below 110.

upvoted 1 times

Pb1805 2 years, 7 months ago

I dont think that R4 will ever take route from R5 since it should be learning internal EIGRP route for the destination route using AD value of 90 and OSPF AD value should be 110.

upvoted 2 times

JOKERR 2 years ago

10.1.1.0/24 is being redistributed into EIGRP on R1 so it has AD of 170.

upvoted 3 times

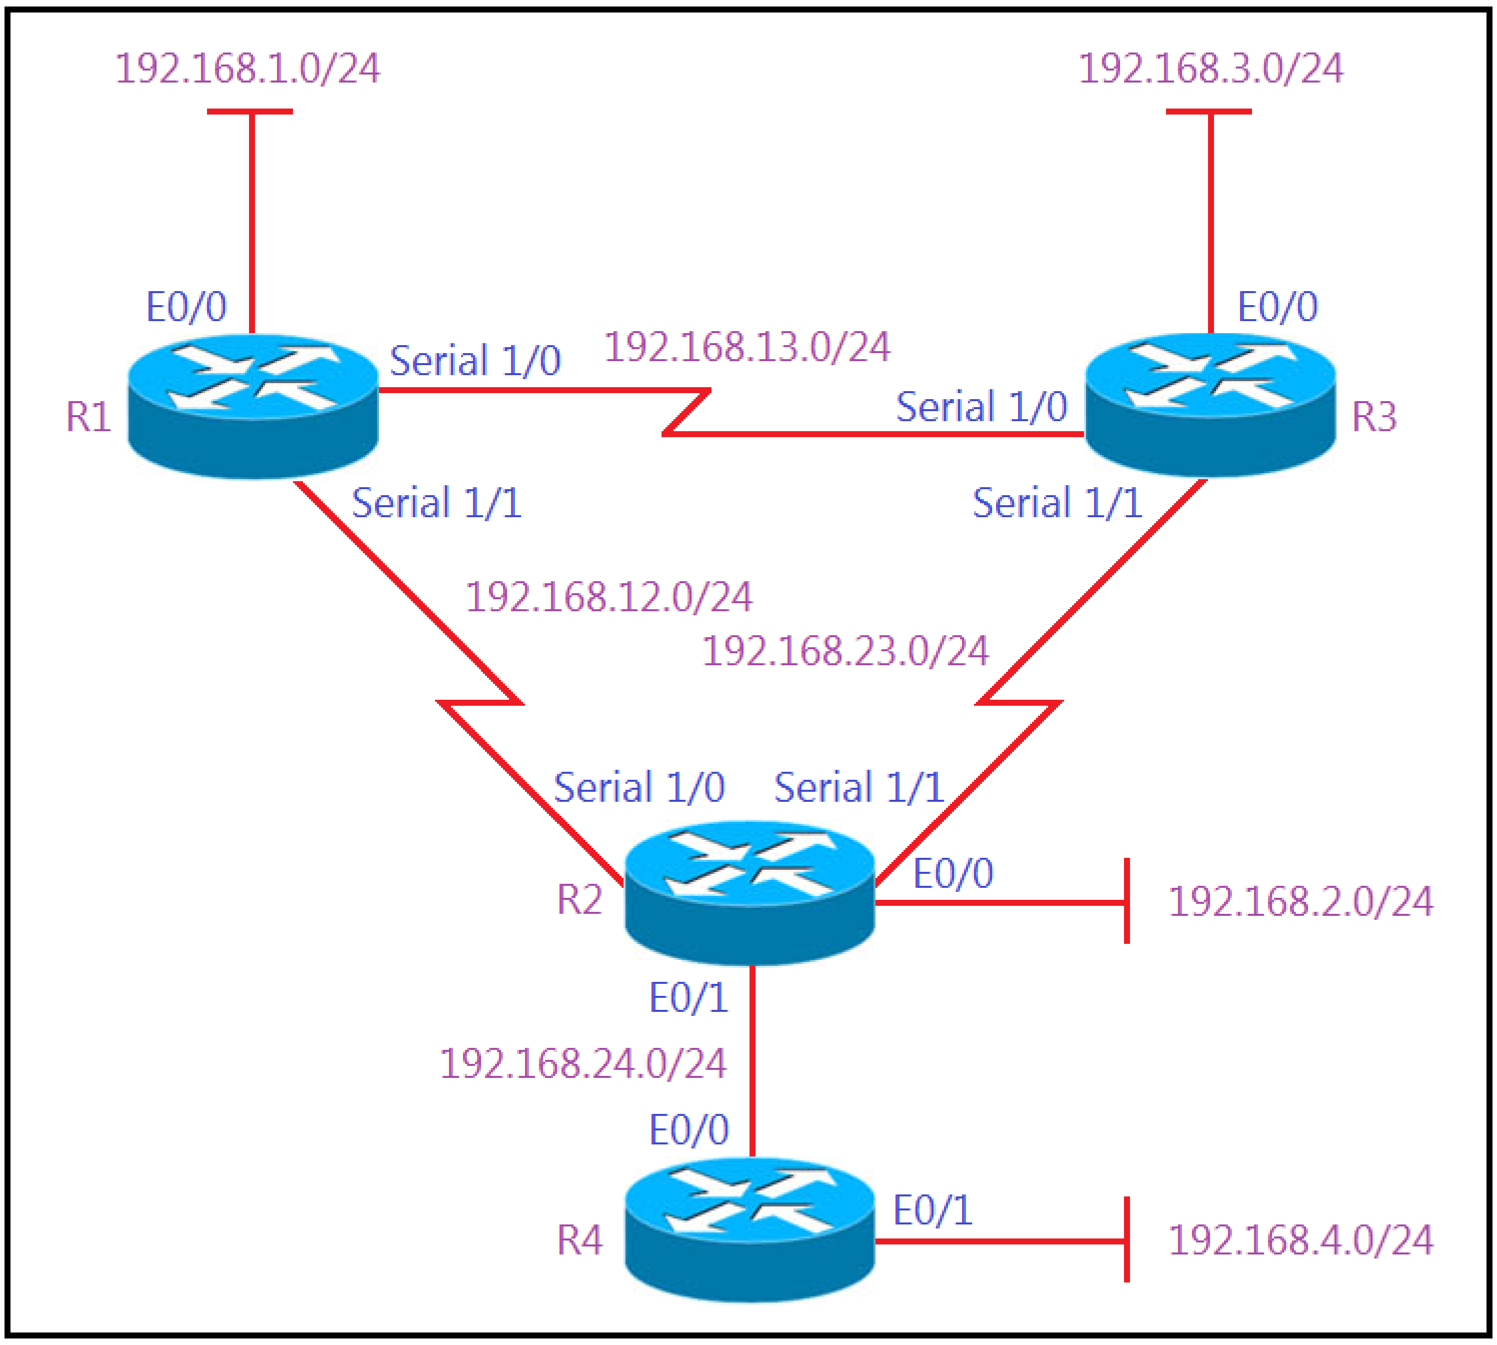

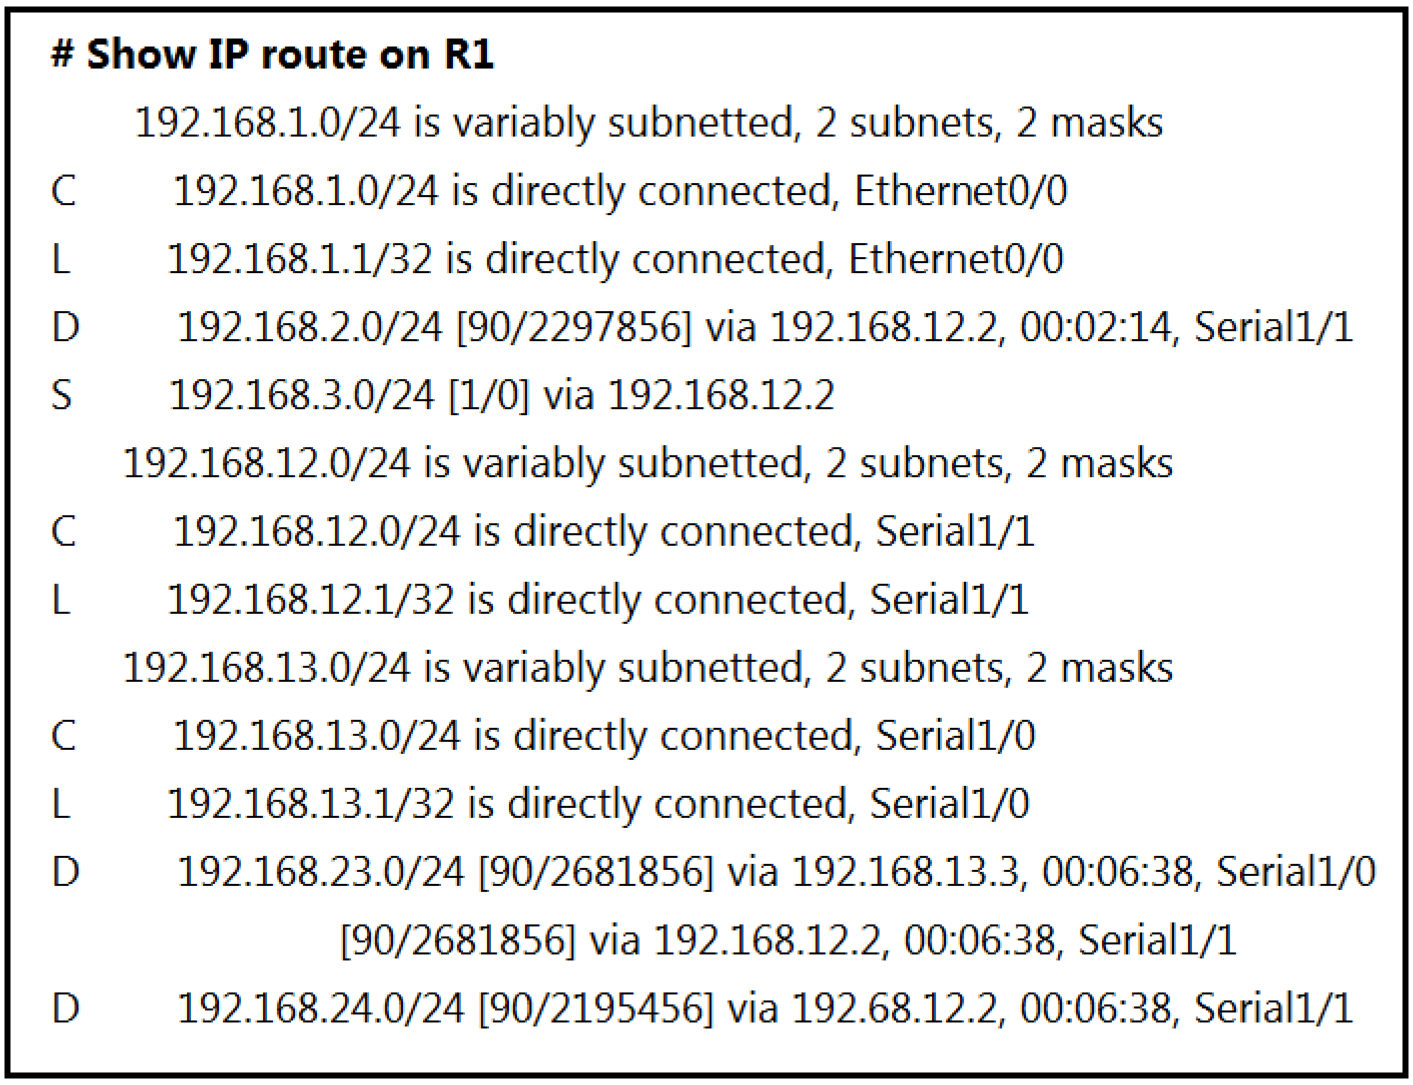

# Show IP route on R1

```
       192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C         192.168.1.0/24 is directly connected, Ethernet0/0
L         192.168.1.1/32 is directly connected, Ethernet0/0
D         192.168.2.0/24 [90/2297856] via 192.168.12.2, 00:02:14, Serial1/1
S         192.168.3.0/24 [1/0] via 192.168.12.2
       192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks
C         192.168.12.0/24 is directly connected, Serial1/1
L         192.168.12.1/32 is directly connected, Serial1/1
       192.168.13.0/24 is variably subnetted, 2 subnets, 2 masks
C         192.168.13.0/24 is directly connected, Serial1/0
L         192.168.13.1/32 is directly connected, Serial1/0
D         192.168.23.0/24 [90/2681856] via 192.168.13.3, 00:06:38, Serial1/0
                          [90/2681856] via 192.168.12.2, 00:06:38, Serial1/1
D         192.168.24.0/24 [90/2195456] via 192.68.12.2, 00:06:38, Serial1/1
```

Refer to the exhibits. All the serial links between R1, R2, and R3 have the same bandwidth. Users on the 192.168.1.0/24 network report slow response times while they access resources on network 192.168.3.0/24. When a traceroute is run on the path, it shows that the packet is getting forwarded via R2 to R3 although the link between R1 and R3 is still up.

What must the network administrator do to fix the slowness?

A. Add a static route on R1 using the next hop of R3.

B. Remove the static route on R1.

C. Change the Administrative Distance of EIGRP to 5.

D. Redistribute the R1 static route to EIGRP.

Correct Answer: *B*

Community vote distribution

B (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

THE ANSWER CORERCT IS B, BECAUSE, HAVE STATIC ROUTE WITH AD "1 lower than Eeigrp "90"

upvoted 1 times

Nhan 1 year, 3 months ago

The given answer is correct, the static route is on the serial interface, also the ad is 1, that why the traffic chose that link to get to the R3, remove the static route will then the R1 will change the routing table and fix the slow link issue

upvoted 1 times

davdtech 1 year, 6 months ago

If you add another static route it will not overwrite the other one. so I think we are assuming that all routers are advertising their networks. so the best choice is to remove the static route.

upvoted 2 times

Hack4 1 year, 10 months ago

YES THE b ANSWER IS CORRECT

upvoted 1 times

Carl1999 1 year, 10 months ago

Selected Answer: B

B is correct.
easy...

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

cakmamail 2 years, 4 months ago

Given answer is correct. After deleting static route, eigrp learned route will be added to routing table

upvoted 1 times

OakA1 2 years, 2 months ago

How do you know that the connected routes are redistributed in EIGRP. Neither you know if 192.168.3.0/24 is participating in EIGRP. For me it's remove static route via R2 and add a new static that points to R3.
Alternatively, information about EIGRP on R3 should be provided

upvoted 2 times

Masashi_O 2 years, 6 months ago

If the network administrator deletes the static route through R2, there will be no route to 192.168.3.0 in the routing table, so I think it is correct to add a static route with R3 as the next hop, is this wrong?

upvoted 2 times

Surfside92 1 year, 11 months ago

You are right. If as most people say the correct answer is b - and the static route to 192.168.3 is removed on R1 - then R1 has no route to that network - from the output its not learned from eigrp. However we have to assume the route to network 192.168.3/24 will show up in the routing table via eigrp when the static route is removed. Also no other answer fully satisfies the solution. If answer A read - add a floating static route to r3 i would be inclined to go for that - but it doesn't !

upvoted 3 times

spapi0390 2 years ago

Wrong since you can not add two static routes from the same lookup ip add. As far as route 192.168.4.0 which is not participating in te EIGRP process but the route towards the R2 yes then it means redistribute connected is configured.

upvoted 1 times

JOKERR 2 years ago

I think given answer is correct. But you CAN add 2 static routes to same destination.

```
ip route 50.0.0.0 255.255.255.0 172.16.45.2
ip route 50.0.0.0 255.255.255.0 10.1.1.2

C1#sh ip rout
S 50.0.0.0 [1/0] via 172.16.45.2
[1/0] via 10.1.1.2
```

upvoted 1 times

⊟ **[Removed]** 1 year, 11 months ago

Only think with configuring 2 statics is you dont know which route it will take and it would still leave users complaining.

upvoted 2 times

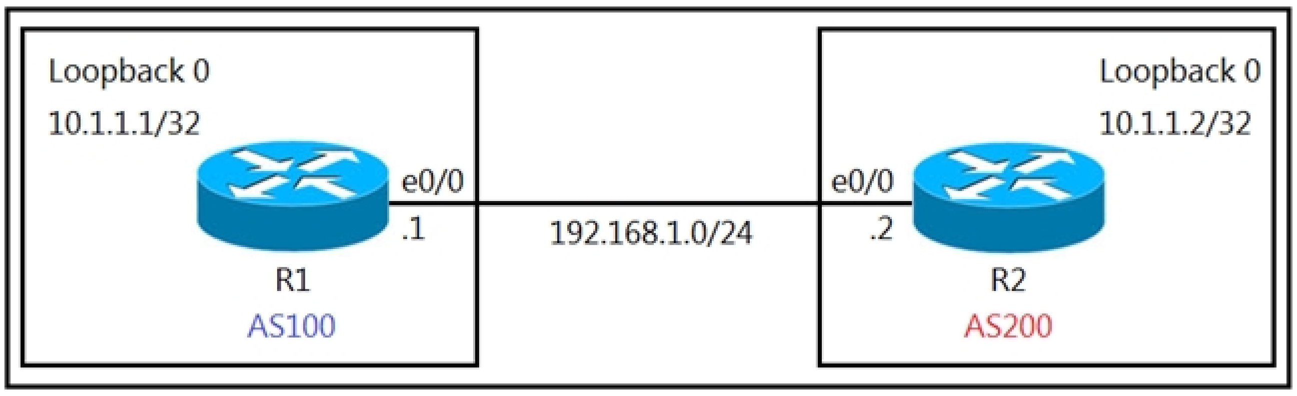

Refer to the exhibit. The R1 and R2 configurations are:

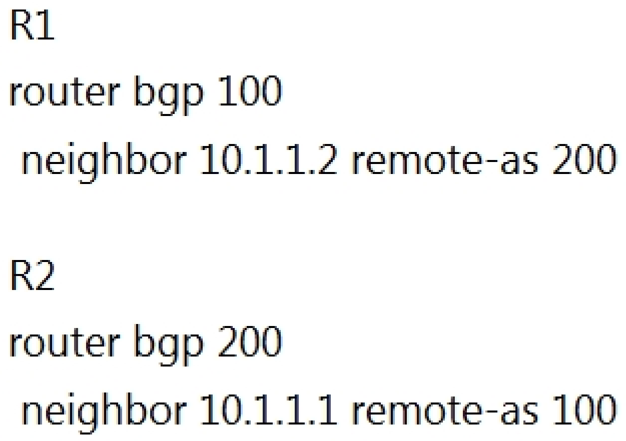

```
R1
router bgp 100
 neighbor 10.1.1.2 remote-as 200

R2
router bgp 200
 neighbor 10.1.1.1 remote-as 100
```

The neighbor relationship is not coming up.

Which two sets of configurations bring the neighbors up? (Choose two.)

A.

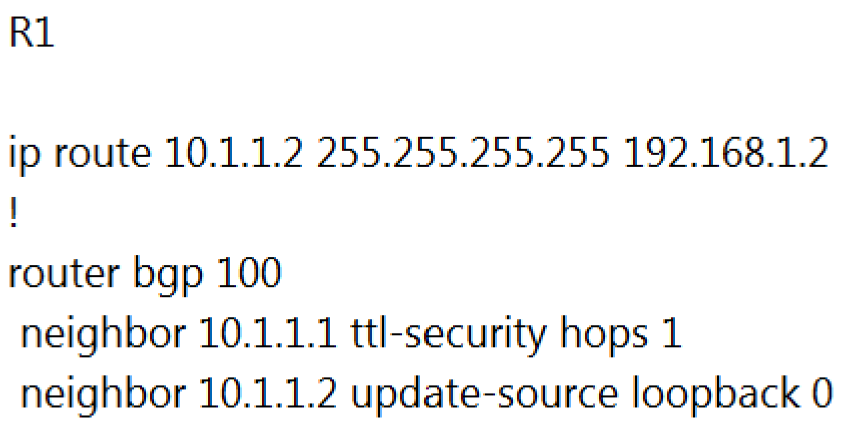

```
R1

ip route 10.1.1.2 255.255.255.255 192.168.1.2
!
router bgp 100
 neighbor 10.1.1.1 ttl-security hops 1
 neighbor 10.1.1.2 update-source loopback 0
```

B.

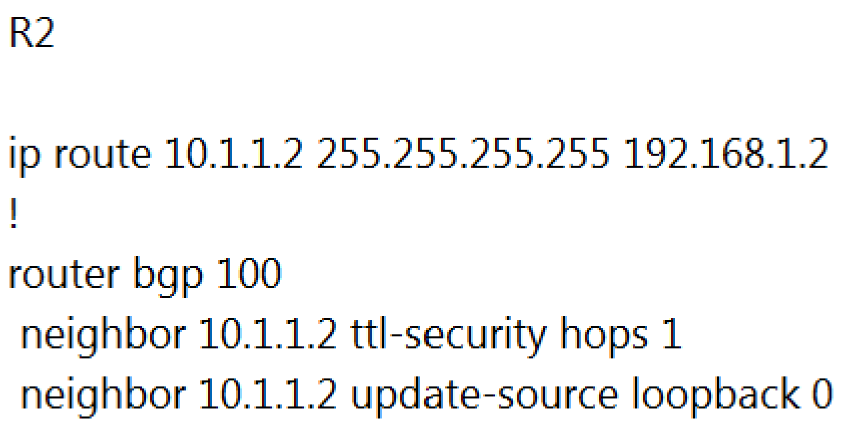

```
R2

ip route 10.1.1.2 255.255.255.255 192.168.1.2
!
router bgp 100
 neighbor 10.1.1.2 ttl-security hops 1
 neighbor 10.1.1.2 update-source loopback 0
```

C.

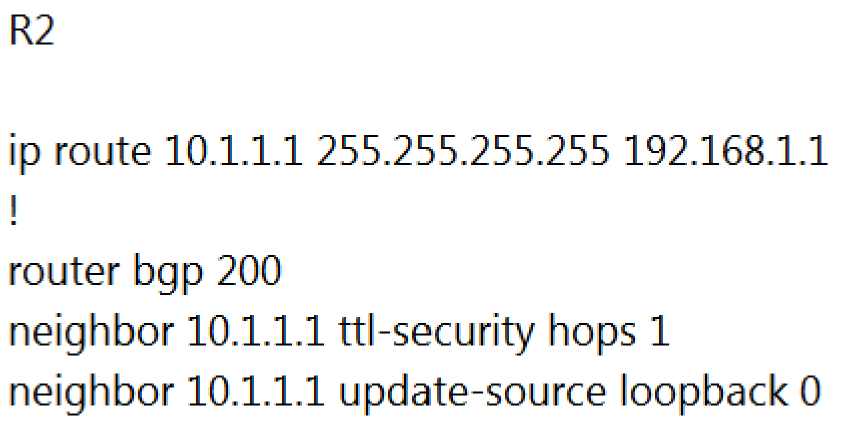

```
R2

ip route 10.1.1.1 255.255.255.255 192.168.1.1
!
router bgp 200
neighbor 10.1.1.1 ttl-security hops 1
neighbor 10.1.1.1 update-source loopback 0
```

D.

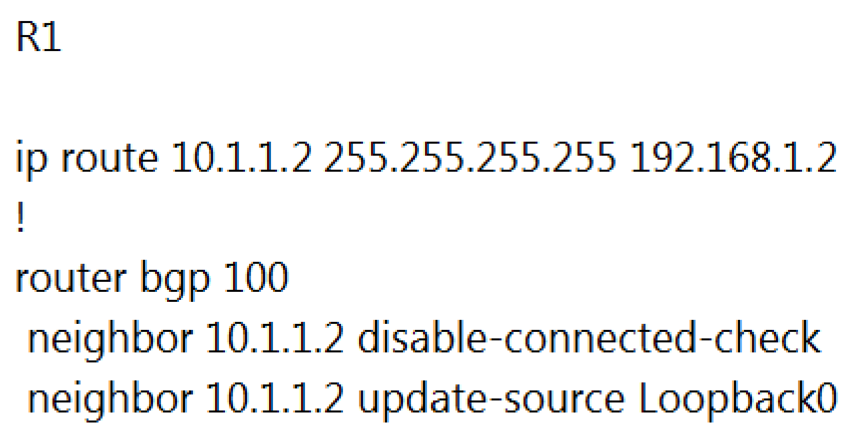

```
R1

ip route 10.1.1.2 255.255.255.255 192.168.1.2
!
router bgp 100
 neighbor 10.1.1.2 disable-connected-check
 neighbor 10.1.1.2 update-source Loopback0
```

E.

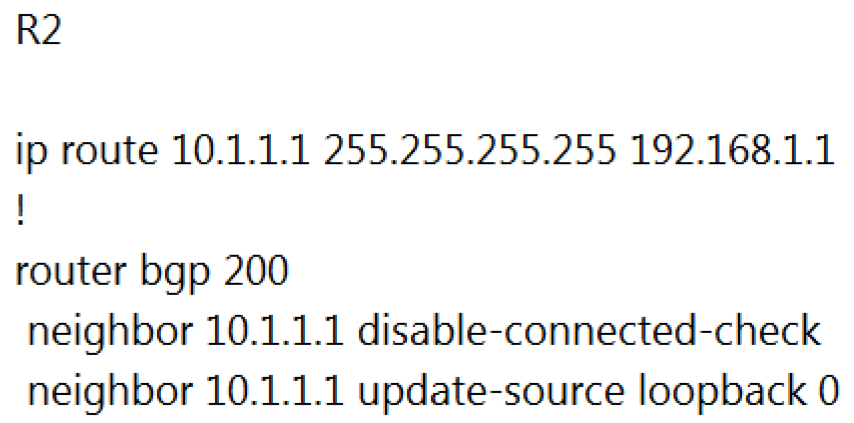

```
R2

ip route 10.1.1.1 255.255.255.255 192.168.1.1
!
router bgp 200
 neighbor 10.1.1.1 disable-connected-check
 neighbor 10.1.1.1 update-source loopback 0
```

Correct Answer: *DE*

conft 4 months ago

D and E is the correct.

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

D and E anwser correct. I test in my lab. Its works

upvoted 1 times

HungarianDish 6 months, 4 weeks ago

Confirmed solution "D"+"E" in CML lab.

upvoted 1 times

drxz 7 months, 3 weeks ago

De and E is correct ;
The disable-connected-check was created precisely for the purpose of peering two directly connected routers on their loopbacks without using the ebgp-multihop
https://ipwithease.com/using-disable-connected-check-in-cisco-bgp/

upvoted 3 times

yonig 8 months, 2 weeks ago

the only problem i have is that the loopbakc is not directly connected - you need another hop. but the ttl security command works the opposite from the multi-hop command. so the value should be 254 or multu hop 2. am i wrong ? unless you have reachabiity using IGP.

upvoted 1 times

Huntkey 1 year, 3 months ago

I didn't lab this but I think "ttl-security hops" have to be at least 2 for it to work. Therefore, the answer is correct.

upvoted 1 times

Hack4 1 year, 10 months ago

The given answer is correct

upvoted 2 times

Carl1999 1 year, 10 months ago

・About the operation of disable-connected-check

When disable-connected-check is set, regardless of the referenced route
Send the packet with TTL = 1.

With disable-connected-check not set,
The router that is trying to send a packet with TTL = 1 has the referenced route
If it is "C" (direct connection), send a packet.
"S" (Statec route) or "O" (OSPF) ,do not send packets.

upvoted 3 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 3 times

Question #54 *Topic 1*

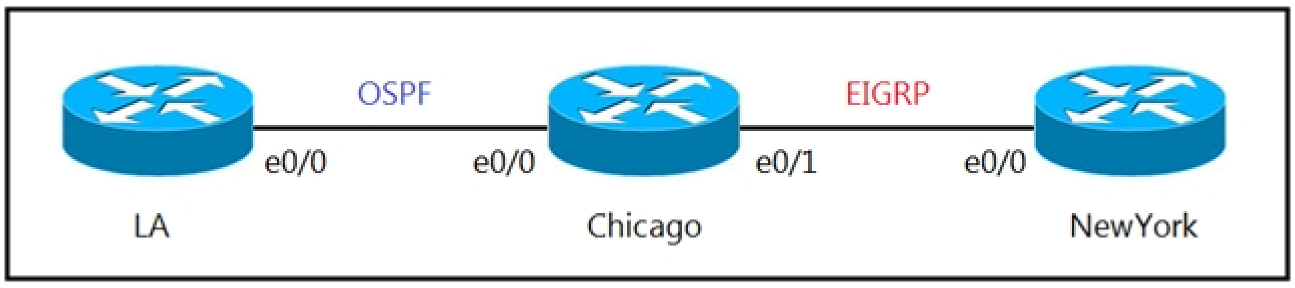

Refer to the exhibit. The network administrator must mutually redistribute routes at the Chicago router to the LA and NewYork routers. The configuration of the
Chicago router is this:

```
router ospf 1
 redistribute eigrp 100
router eigrp 100
 redistribute ospf 1
```

After the configuration, the LA router receives all the NewYork routes, but the NewYork router does not receive any LA routes.

Which set of configurations fixes the problem on the Chicago router?

A.

```
router ospf 1
    redistribute eigrp 100 metric 20
```

B.

```
router eigrp 100
    redistribute ospf 1 metric 10 10 10 10 10
```

C.

```
router ospf 1
    redistribute eigrp 100 subnets
```

D.

```
router eigrp 100
    redistribute ospf 1 subnets
```

Correct Answer: *B*

robi1020 9 months ago

We have to specify a metric, if we don't, redistribution fails.

EIGRP and OSPF use different metrics and there is no way to convert from one metric to another. This means we have to configure the metric ourselves.

EIGRP uses a metric that is based on bandwidth, delay, reliability, load, and MTU (even though MTU is not actually used in the calculation).

upvoted 3 times

Hurk2 11 months, 2 weeks ago

This question should have chose two, redistribute eigrp 1 subnets is also needed for O E2 routes to populate in LA

upvoted 2 times

timtgh 1 year, 6 months ago

Wouldn't they have gotten an error message if they typed the command without the metrics?

upvoted 1 times

JOKERR 1 year, 6 months ago

There would be no error message if the redistribution command is typed without metrics. In that case, metric would be set to Infinity (unreachable).

upvoted 4 times

Hack4 1 year, 10 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct
upvoted 2 times

Question #55

Topic 1

DRAG DROP -

Drag and drop the actions from the left into the correct order on the right to configure a policy to avoid following packet forwarding based on the normal routing path.

Select and Place:

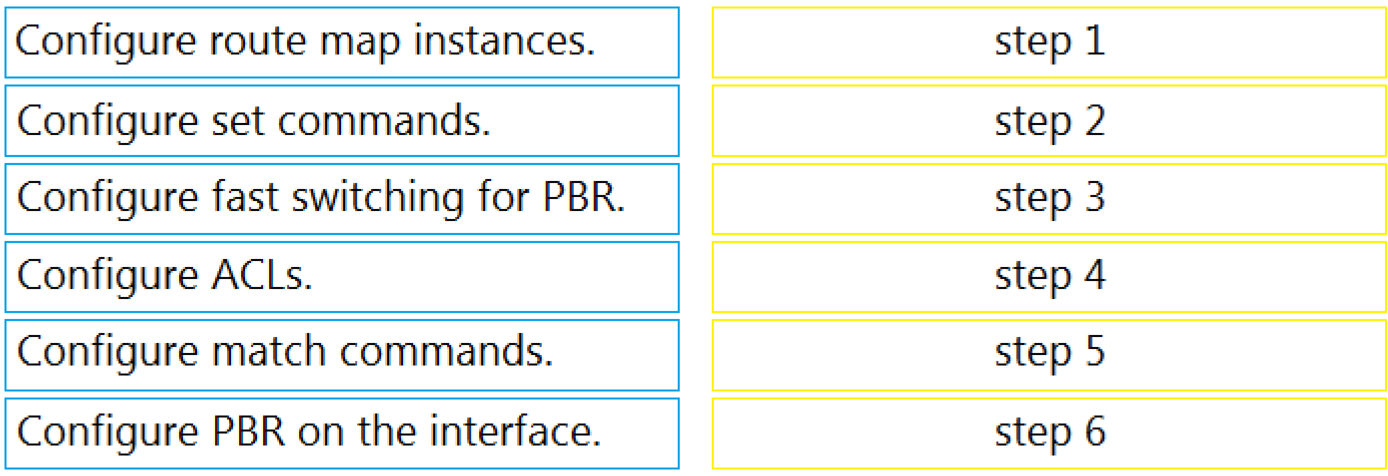

Configure route map instances.	step 1
Configure set commands.	step 2
Configure fast switching for PBR.	step 3
Configure ACLs.	step 4
Configure match commands.	step 5
Configure PBR on the interface.	step 6

Correct Answer:

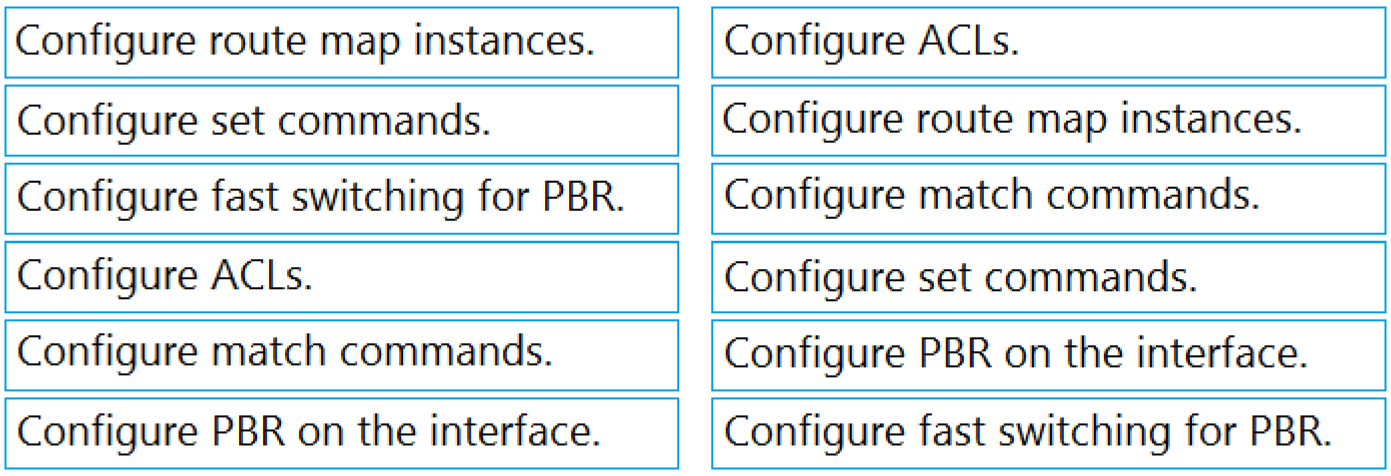

Configure route map instances.	Configure ACLs.
Configure set commands.	Configure route map instances.
Configure fast switching for PBR.	Configure match commands.
Configure ACLs.	Configure set commands.
Configure match commands.	Configure PBR on the interface.
Configure PBR on the interface.	Configure fast switching for PBR.

Reference:

https://community.cisco.com/t5/networking-documents/how-to-configure-pbr/ta-p/3122774

conft 4 months ago

the given answer is correct.
upvoted 1 times

JOKERR 1 year, 6 months ago

Give answer is correct.

Here is a better source:

https://howdoesinternetwork.com/2013/configuration-of-pbr-policy-based-routing
upvoted 2 times

Vainius 2 years, 2 months ago

https://community.cisco.com/kxiwq67737/attachments/kxiwq67737/6016-discussions-lan-switching-routing/26658/1/8609-PBR%20configuration.pdf
upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct
upvoted 1 times

Question #56 *Topic 1*

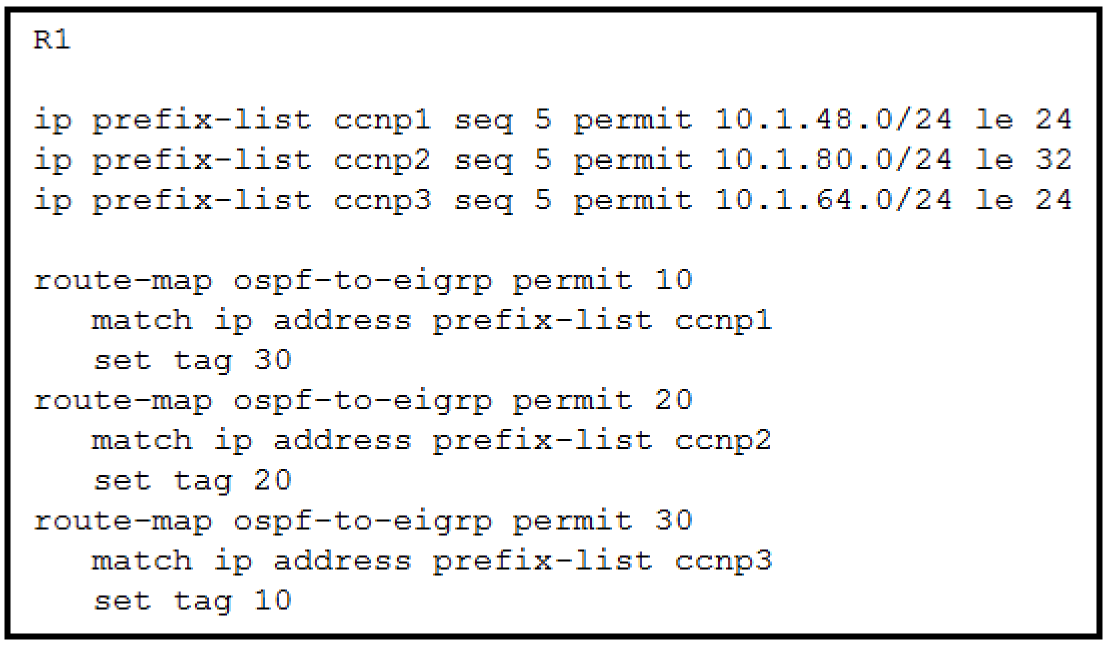

```
R1

ip prefix-list ccnp1 seq 5 permit 10.1.48.0/24 le 24
ip prefix-list ccnp2 seq 5 permit 10.1.80.0/24 le 32
ip prefix-list ccnp3 seq 5 permit 10.1.64.0/24 le 24

route-map ospf-to-eigrp permit 10
   match ip address prefix-list ccnp1
   set tag 30
route-map ospf-to-eigrp permit 20
   match ip address prefix-list ccnp2
   set tag 20
route-map ospf-to-eigrp permit 30
   match ip address prefix-list ccnp3
   set tag 10
```

Refer to the exhibit. An engineer wanted to set a tag of 30 to route 10.1.80.65/32 but it failed.
How is the issue fixed?

A. Modify route-map ospf-to-eigrp permit10 and match prefix-list ccnp2.

B. Modify prefix-list ccnp3 to add 10.1.64.0/20 ge 32.

C. Modify prefix-list ccnp3 to add 10.1.64.0/20 le 24.

D. Modify route-map ospf-to-eigrp permit 30 and match prefix-list ccnp2.

Correct Answer: *D*

Community vote distribution

A (100%)

Dave22 Highly Voted 2 years, 7 months ago
I chose A as D just does not make sense it would set a tag to be 10 not 30
upvoted 15 times

RHK0783 Highly Voted 2 years, 7 months ago
A is correct ...
upvoted 9 times

Chiaretta Most Recent 5 months ago
Selected Answer: A
A is the right answer.
upvoted 1 times

Dacusai 8 months ago
A is the correct one
upvoted 1 times

Hurk2 11 months, 2 weeks ago
Selected Answer: A
A is correct
upvoted 1 times

baldebri 11 months, 4 weeks ago
D is the correct one, the prefix list is always added to the end unless the sequence keyword is mentioned so to make any changes modify the last or add seq 40
upvoted 1 times

mrnipsnips 1 year, 1 month ago
Selected Answer: A
A is correct D doesnt make any sense
upvoted 1 times

Alexloh 1 year, 5 months ago
Selected Answer: A
A is the correct answer

upvoted 1 times

Bronco30A 1 year, 7 months ago

Selected Answer: A

A is correct

upvoted 2 times

Alex147 1 year, 9 months ago

Selected Answer: A

A is correct.

upvoted 1 times

Hack4 1 year, 10 months ago

A is the best answer

upvoted 1 times

wts 1 year, 10 months ago

Selected Answer: A

It is necessary that "ccnp2" and "30" are in the same route map.

upvoted 1 times

Girmiti 1 year, 11 months ago

Selected Answer: A

correct answer is A. 10.1.80.65/32 needed to add tag 30 so we must modify route-map ospf-to-eigrp permit 10 to match ip prefix-list ccnp2

upvoted 1 times

Chris_Li 2 years ago

Selected Answer: A

i think the answer is A

upvoted 3 times

examShark 2 years, 4 months ago

The correct answer is A

upvoted 6 times

Pb1805 2 years, 7 months ago

This seems straight forward....Answer is A

upvoted 4 times

ZachTL11 2 years, 8 months ago

Correct Answer:
Modify route-map ospf-to-eigrp permit10 and match prefix-list ccnp2

upvoted 1 times

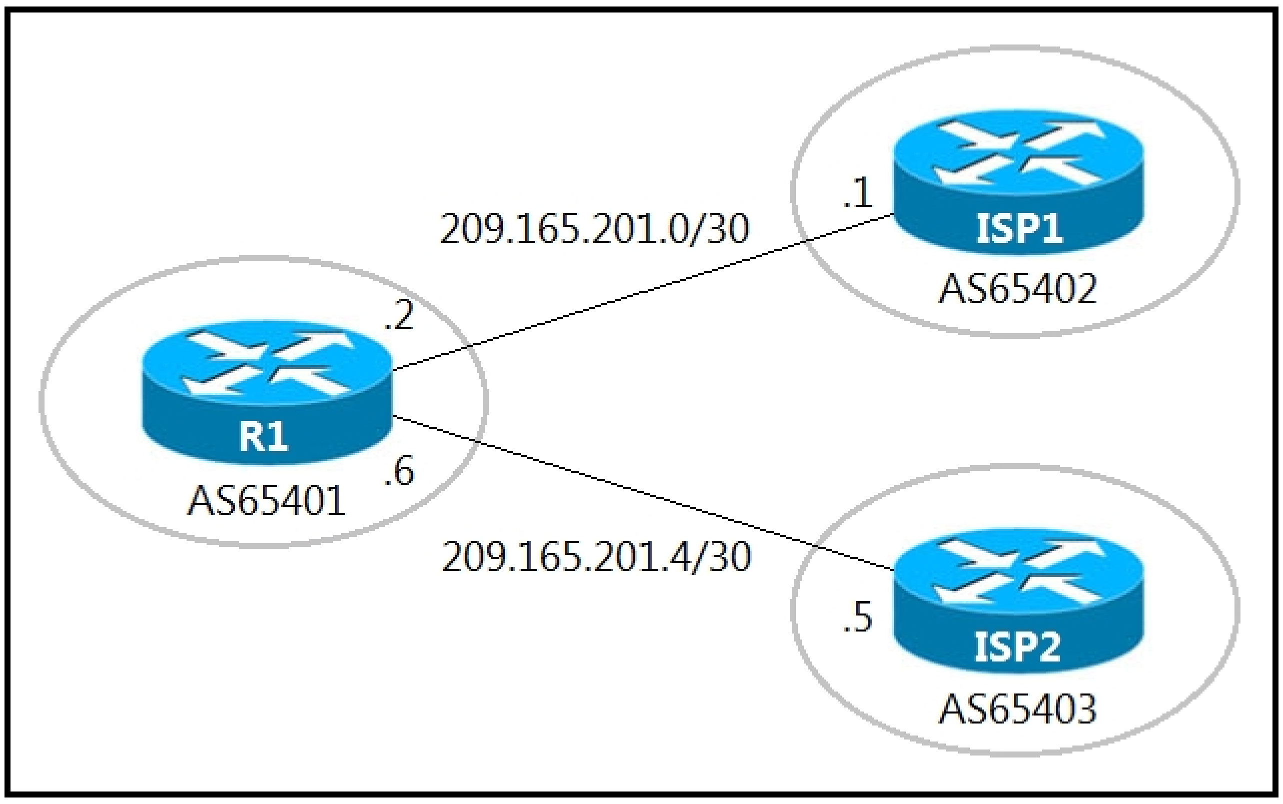

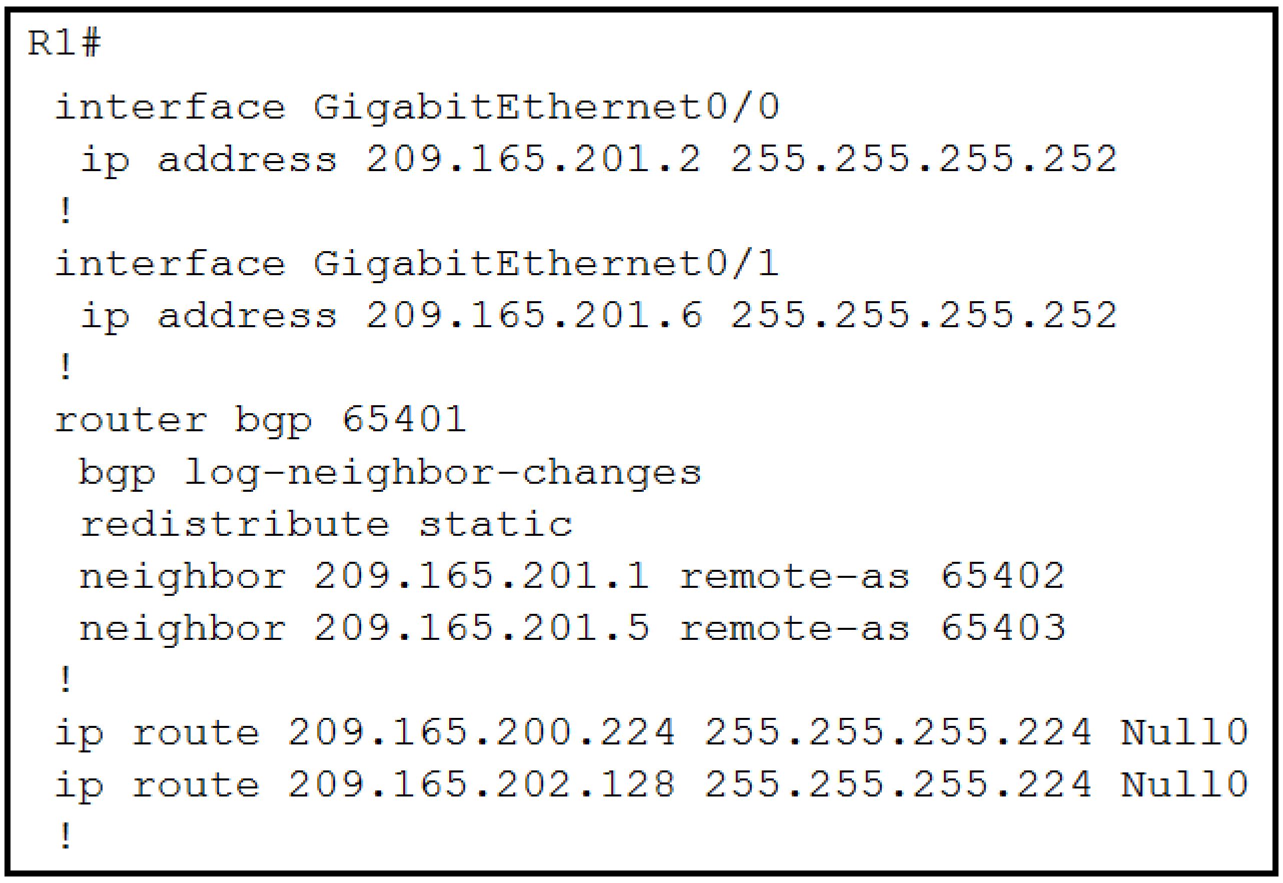

```
R1#
 interface GigabitEthernet0/0
  ip address 209.165.201.2 255.255.255.252
 !
 interface GigabitEthernet0/1
  ip address 209.165.201.6 255.255.255.252
 !
 router bgp 65401
  bgp log-neighbor-changes
  redistribute static
  neighbor 209.165.201.1 remote-as 65402
  neighbor 209.165.201.5 remote-as 65403
 !
 ip route 209.165.200.224 255.255.255.224 Null0
 ip route 209.165.202.128 255.255.255.224 Null0
 !
```

Refer to the exhibits. A company with autonomous system number AS65401 has obtained IP address block 209.165.200.224/27 from ARIN. The company needed more IP addresses and was assigned block 209.165.202.128/27 from ISP2. An engineer in ISP1 reports that they are receiving ISP2 routes from AS65401.

Which configuration on R1 resolves the issue?

A.

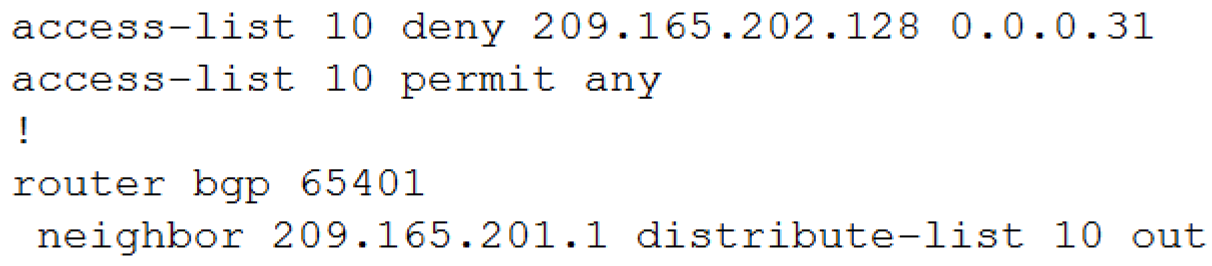

```
access-list 10 deny 209.165.202.128 0.0.0.31
access-list 10 permit any
 !
router bgp 65401
 neighbor 209.165.201.1 distribute-list 10 out
```

B.

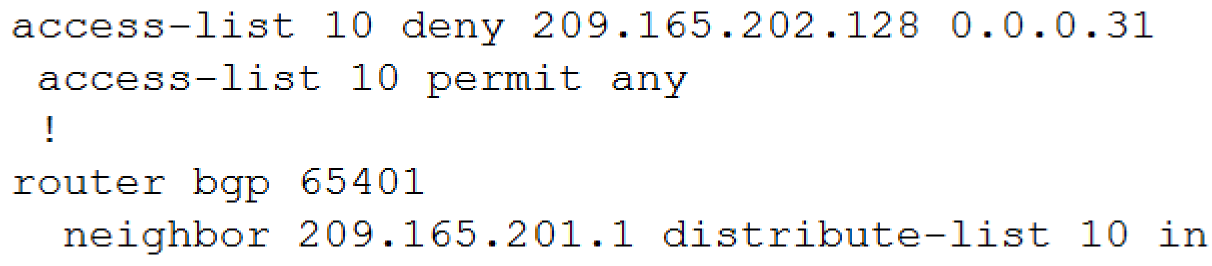

```
access-list 10 deny 209.165.202.128 0.0.0.31
 access-list 10 permit any
  !
router bgp 65401
  neighbor 209.165.201.1 distribute-list 10 in
```

C.

```
ip route 209.165.200.224 255.255.255.224 209.165.201.1
ip route 209.165.202.128 255.255.255.224 209.165.201.5
```

D.

```
ip route 0.0.0.0 0.0.0.0 209.165.201.1
ip route 0.0.0.0 0.0.0.0 100 209.165.201.5
```

Correct Answer: *A*

Nhan 1 year, 3 months ago

The given answer is correct, simple acl block the route coming in and redistribute to the ISP 1 from us

upvoted 1 times

examShark 2 years, 4 months ago

Te given answer is correct

upvoted 3 times

RTE 2 years, 5 months ago

in answer invalid neighbour statement, but direction and access list are good

upvoted 1 times

cakmamail 2 years, 4 months ago

Ulan burada da mı çıktın karşıma!

upvoted 2 times

Question #58 *Topic 1*

After some changes in the routing policy, it is noticed that the router in AS 45123 is being used as a transit AS router for several service providers.

Which configuration ensures that the branch router in AS 45123 advertises only the local networks to all SP neighbors?

A.

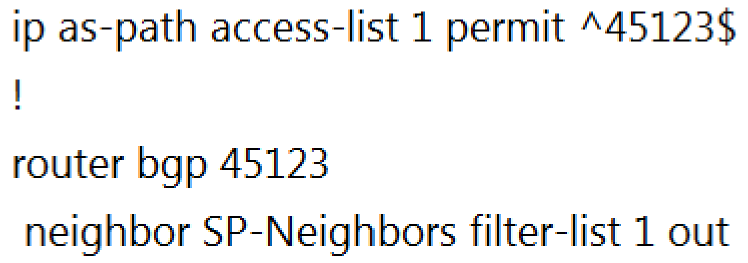

```
ip as-path access-list 1 permit ^45123$
!
router bgp 45123
 neighbor SP-Neighbors filter-list 1 out
```

B.

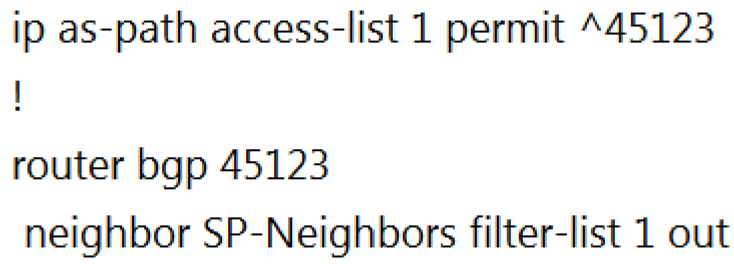

```
ip as-path access-list 1 permit ^45123
!
router bgp 45123
 neighbor SP-Neighbors filter-list 1 out
```

C.

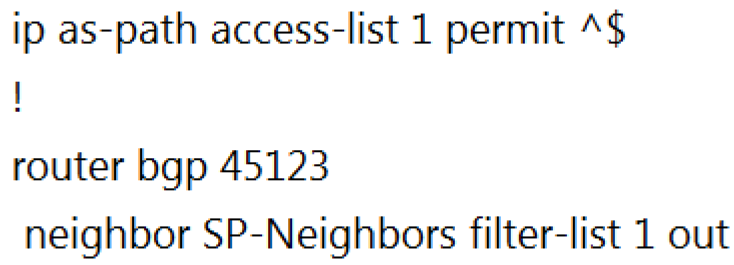

```
ip as-path access-list 1 permit ^$
!
router bgp 45123
 neighbor SP-Neighbors filter-list 1 out
```

D.

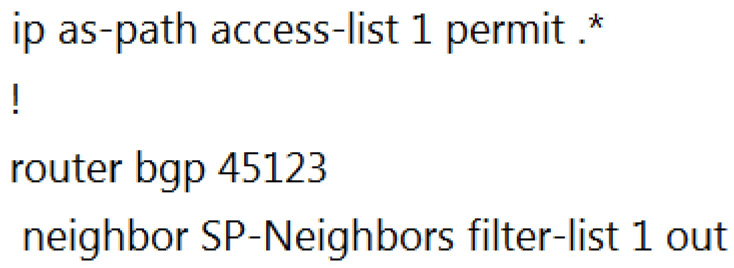

```
ip as-path access-list 1 permit .*
!
router bgp 45123
 neighbor SP-Neighbors filter-list 1 out
```

Correct Answer: *C*

⊟ **Alexloh** Highly Voted 1 year, 5 months ago

the regular expression (^$) matches any route that has an empty AS path attribute (that is, no character from start to end). Only locally originated routes have an empty AS path attribute; hence this regular expression is used when matching local routes. This type of filter is used by multihomed customers to send only their address space to their service providers, to prevent them from becoming a transit AS.

upvoted 11 times

⊟ **Mjestic** Highly Voted 2 years, 3 months ago

C is correct.
-> https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/26634-bgp-toc.html#asregexp
--> ^$: This expression indicates origination from this AS.

upvoted 10 times

⊟ **conft** Most Recent 4 months ago

C is correct.

upvoted 1 times

⊟ **Hack4** 1 year, 10 months ago

C is correct

upvoted 1 times

⊟ **Dirkd0344** 2 years ago

It is C. The ^$ regex statement matches an empty AS Path. Only locally originated routes will have an empty AS Path.

upvoted 2 times

⊟ **examShark** 2 years, 4 months ago

The correct answer is A
^$ ^ = begins with, $ = ends with. What BGP routes have no ASN!

upvoted 1 times

mosvan 2 years, 4 months ago

@examShark, Thank you for your contribution, but here you are wrong.
Paths originating locally can be matched by ^$ and filter out to ISPs.
So answer C is the correct one.

upvoted 3 times

vdsdrs 2 years, 4 months ago

While a newly sourced route is still within the AS in which it was created, the AS path is empty. When the AS has a requirement to filter out all but the routes that are local to itself before sending them to a neighboring AS, the AS will permit sending of the routes with the empty AS path and will deny all others.
Answer C

upvoted 2 times

RTE 2 years, 5 months ago

It's B, locally orginitated - started with ASN of enterpise and ends on this ASN.
Other uses match characterf at the end of matching string -$
Correct me

upvoted 1 times

Masashi_O 2 years, 6 months ago

C.
Specify only the routes generated by your own AS.

upvoted 2 times

Question #59

Topic 1

A network administrator is troubleshooting a high utilization issue on the route processor of a router that was reported by NMS. The administrator logged into the router to check the control plane policing and observed that the BGP process is dropping a high number of routing packets and causing thousands of routes to recalculate frequently.
Which solution resolves this issue?

A. Shape the pir for BGP, conform-action set-prec-transmit, and exceed action set-frde-transmit.

B. Police the pir for BGP, conform-action set-prec-transmit, and exceed action set-clp-transmit.

C. Shape the cir for BGP, conform-action transmit, and exceed action transmit.

D. Police the cir for BGP, conform-action transmit, and exceed action transmit.

Correct Answer: *D*
Reference:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_plcshp/configuration/xe-3s/qos-plcshp-xe-3s-book/qos-plcshp-plcr-mact.html

Community vote distribution

D (100%)

ZamanR 5 days ago

I think A

Explanation

CIR (Committed Information Rate) is the minimum guaranteed traffic delivered in the network.

PIR (Peak Information Rate) is the top bandwidth point of allowed traffic in a non busy times without any guarantee.

upvoted 1 times

ZamanR 5 days ago

Policing: is used to control the rate of traffic flowing across an interface. During a bandwidth exceed (crossed the maximum configured rate), the excess traffic is generally dropped or remarked. The result of traffic policing is an output rate that appears as a saw-tooth with crests and troughs. Traffic policing can be applied to inbound and outbound interfaces. Unlike traffic shaping, QoS policing avoids delays due to queuing. Policing is configured in bytes.

+ Shaping: retains excess packets in a queue and then schedules the excess for later transmission over increments of time. When traffic reaches the maximum configured rate, additional packets are queued instead of being dropped to proceed later. Traffic shaping is applicable only on outbound interfaces as buffering and queuing happens only on outbound interfaces. Shaping is configured in bits per second.

upvoted 1 times

ZamanR 5 days ago

Therefore in this case we can only policing, not shaping as traffic shaping is applicable only on outbound interfaces as buffering and queuing happens only on outbound interfaces. Moreover, BGP traffic is not important so we can drop the excess packets without any problems.

And we only policing the PIR traffic so that the route processor is not overwhelmed by BGP calculation
Note: The “set-prec-transmit” is the same as “transmit” command except it sets the IP Precedence level as well. The “set-clp-transmit” sets the ATM Cell Loss Priority (CLP) bit from 0 to 1 on the ATM cell and transmits the packet.

upvoted 1 times

AinsB 7 months, 1 week ago

Selected Answer: D

One very important concept is that Traffic shaping allows you to control the speed of traffic that is leaving an interface. This way, you can match the flow of the traffic to the speed of the interface receiving the packet.
Policing works in both directions Input and Output. "Traffic policing allows you to control the maximum rate of traffic sent or received on an interface and to partition a network into multiple priority levels or class of service (CoS)."

upvoted 3 times

HungarianDish 8 months ago

Selected Answer: D

This is how I see it. As Huntkey mentioned, the issue is not described as relating to ATM or Frame Relay, so we can ignore A and B. Then we we need to choose between C and D. C is for shaping, but you won't achieve shaping with the commands "conform-action transmit action-transmit", so C is not a valid solution. Excluding wrong answers, D is left.

upvoted 2 times

juliop 8 months, 2 weeks ago

Why not Police PIR?

upvoted 1 times

MasterMatt 8 months, 2 weeks ago

Policy map control-plane does support both policing and shaping. I'm unsure which one is the correct answer between C and D.

upvoted 1 times

chris7890 1 year, 1 month ago

can someone explain in more detail why the answer is correct?

upvoted 1 times

Huntkey 1 year, 2 months ago

set-clp-transmit set atm clp and send it
set-frde-transmit set FR DE and send it
I guess this is not an ATM or FR circuit

upvoted 1 times

Huntkey 1 year, 3 months ago

I think it is called control plane POLICING and not SHAPING is because it only supports policing and not shaping. So D is correct

upvoted 3 times

Audie 1 year, 9 months ago

I think C...Shape the cir for BGP....in order to reduce BGP recalculation

upvoted 2 times

Carl1999 1 year, 10 months ago

given answer is correct.
It needs policing the CIR.

upvoted 1 times

Question #60 *Topic 1*

Which mechanism must be chosen to optimize the reconvergence time for OSPF at company location 408817202 that is less CPU-intensive than reducing the hello and dead timers?

A. sso

B. BFD

C. Dead Peer Detection keepalives

D. OSPF demand circuit

Correct Answer: *B*

Reference:

https://forum.networklessons.com/t/ospf-hello-and-dead-interval/1255

Alexloh 1 year, 5 months ago

B is the best answer, the rest of the answers look inrelevant.

upvoted 2 times

Hack4 1 year, 10 months ago

B is correct

upvoted 1 times

doumba 1 year, 11 months ago

the given answer is correct

upvoted 1 times

Question #61

Topic 1

Refer to the exhibit.

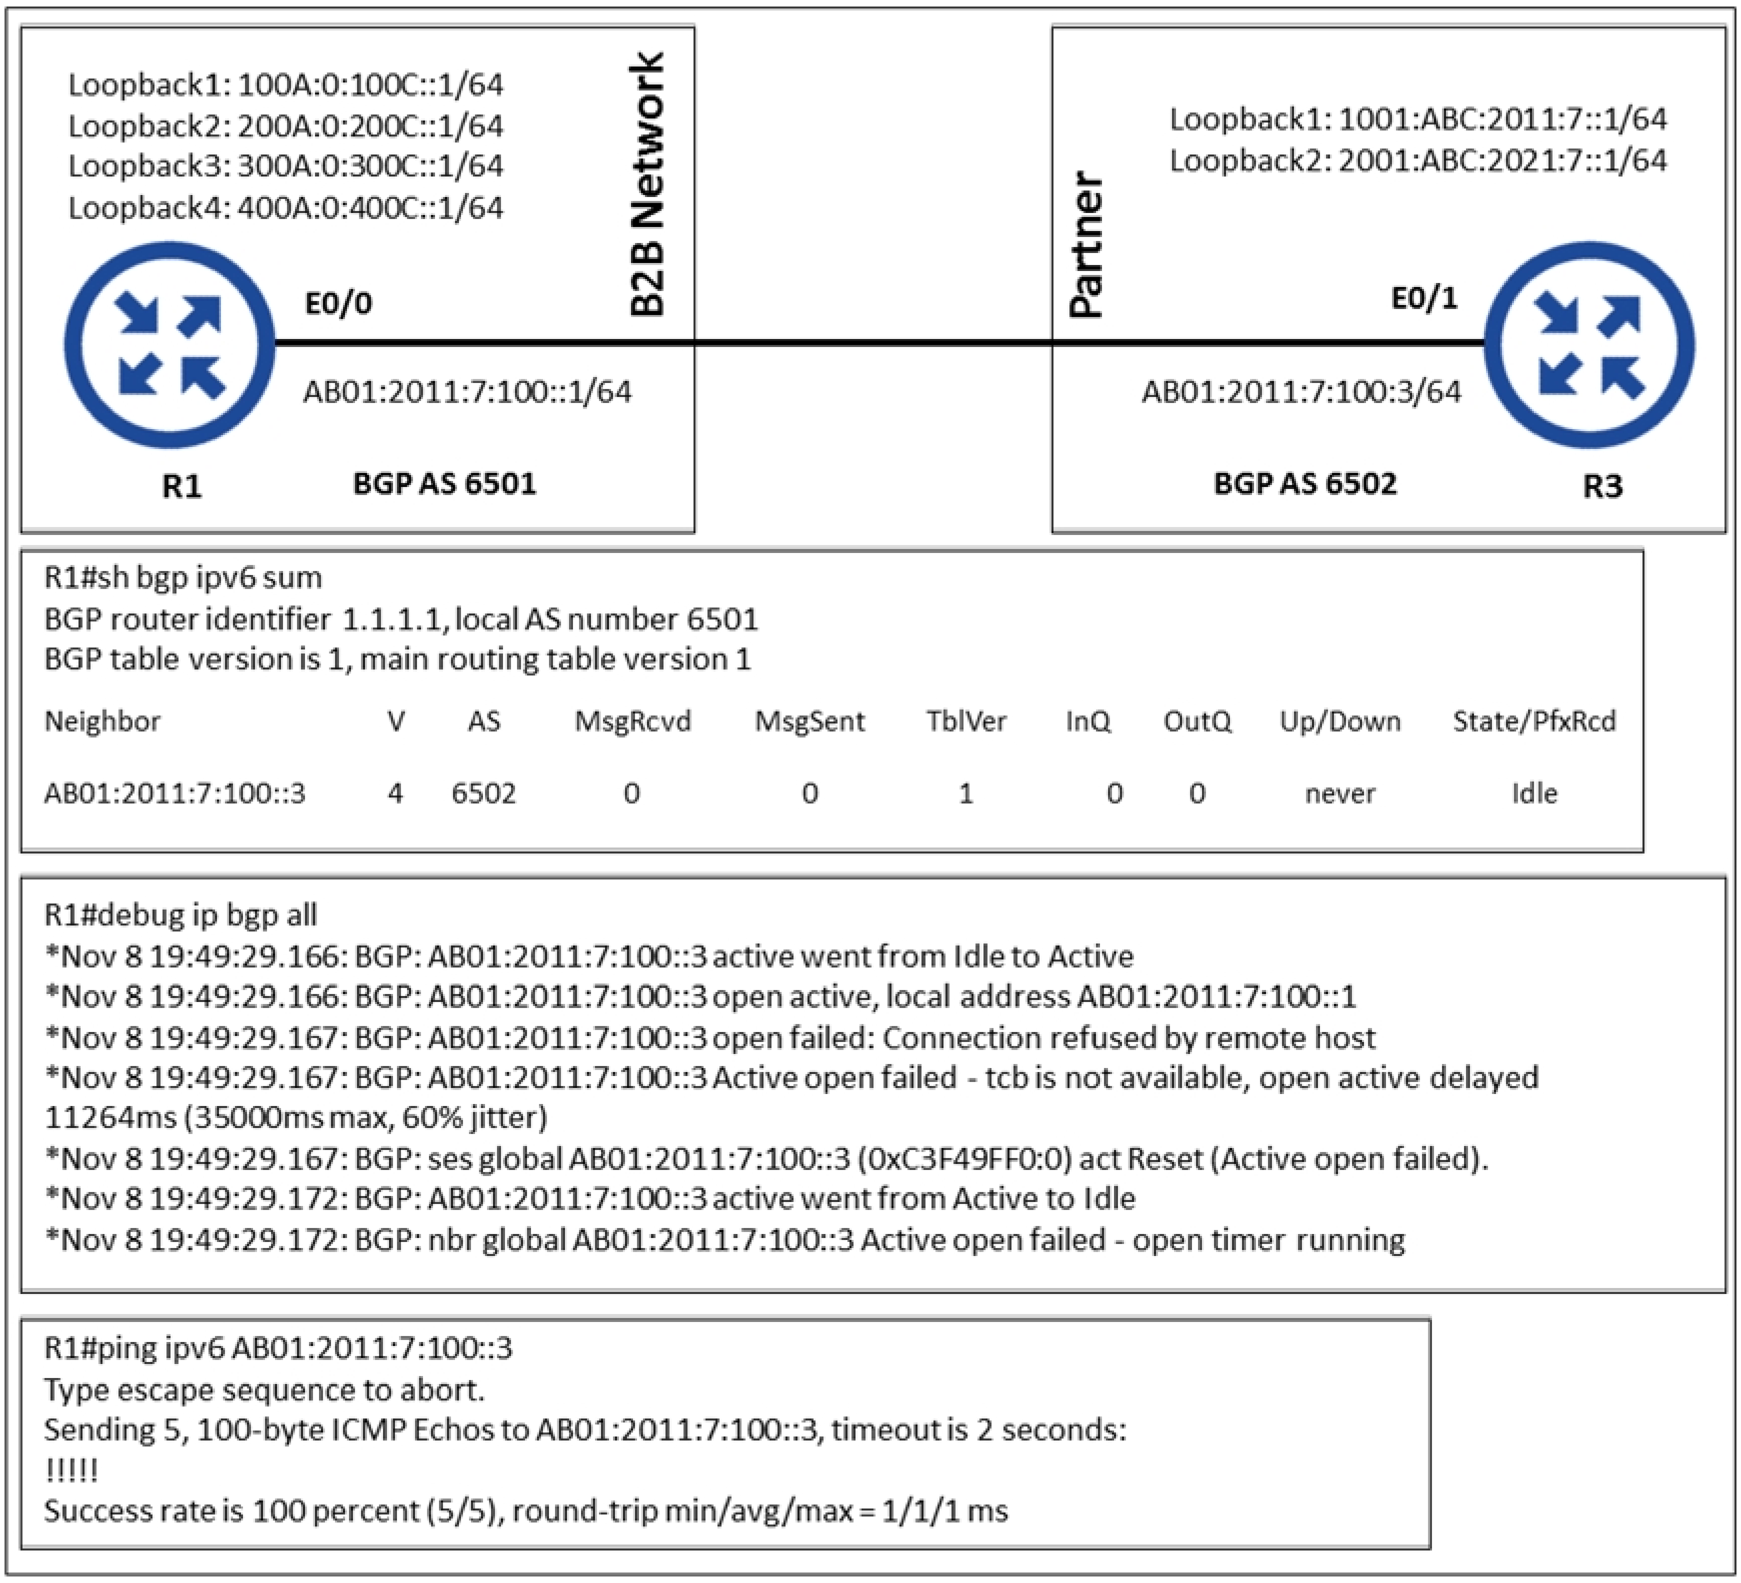
Loopback1: 100A:0:100C::1/64
Loopback2: 200A:0:200C::1/64
Loopback3: 300A:0:300C::1/64
Loopback4: 400A:0:400C::1/64
B2B Network
E0/0
AB01:2011:7:100::1/64
R1
BGP AS 6501
Partner
Loopback1: 1001:ABC:2011:7::1/64
Loopback2: 2001:ABC:2021:7::1/64
E0/1
AB01:2011:7:100:3/64
BGP AS 6502
R3

```
R1#sh bgp ipv6 sum
BGP router identifier 1.1.1.1, local AS number 6501
BGP table version is 1, main routing table version 1
Neighbor              V    AS     MsgRcvd    MsgSent    TblVer   InQ   OutQ   Up/Down    State/PfxRcd
AB01:2011:7:100::3    4    6502      0          0          1       0     0      never        Idle
```

```
R1#debug ip bgp all
*Nov 8 19:49:29.166: BGP: AB01:2011:7:100::3 active went from Idle to Active
*Nov 8 19:49:29.166: BGP: AB01:2011:7:100::3 open active, local address AB01:2011:7:100::1
*Nov 8 19:49:29.167: BGP: AB01:2011:7:100::3 open failed: Connection refused by remote host
*Nov 8 19:49:29.167: BGP: AB01:2011:7:100::3 Active open failed - tcb is not available, open active delayed
11264ms (35000ms max, 60% jitter)
*Nov 8 19:49:29.167: BGP: ses global AB01:2011:7:100::3 (0xC3F49FF0:0) act Reset (Active open failed).
*Nov 8 19:49:29.172: BGP: AB01:2011:7:100::3 active went from Active to Idle
*Nov 8 19:49:29.172: BGP: nbr global AB01:2011:7:100::3 Active open failed - open timer running
```

```
R1#ping ipv6 AB01:2011:7:100::3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to AB01:2011:7:100::3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

An engineer configured BGP between routers R1 and R3. The BGP peers cannot establish neighbor adjacency to be able to exchange routes. Which configuration resolves this issue?

A. R1 router bgp 6501 address-family ipv6 neighbor AB01:2011:7:100::3 activate

B. R3 router bgp 6502 address-family ipv6 neighbor AB01:2011:7:100::1 activate

C. R1 router bgp 6501 neighbor AB01:2011:7:100::3 ebgp-multihop 255

D. R3 router bgp 6502 neighbor AB01:2011:7:100::1 ebgp-multihop 255

Correct Answer: *B*

Community vote distribution

B (100%)

networkWiz Highly Voted 1 year, 4 months ago

Selected Answer: B

B is the correct answer.

As it states in the debug "Connection refused by remote host".
Extra step needed on the remote router (R3) which is to activate the neighbor in address-family ipv6 unicast and run the "neighbor <neighbor_IP> activate" command.

upvoted 5 times

Nhan Most Recent 1 year, 3 months ago

The given answer is correct, because the pong show that the R3 responded, which meant the configuration is correct, then we must look at R1 to see what chase the issue

upvoted 1 times

⊟ **leogp79** 1 year, 4 months ago

I just testes this scenrio on GNS3, and B is the correct answer

upvoted 1 times

⊟ **Reikidude00** 1 year, 5 months ago

Selected Answer: B

It's B 4 sure

upvoted 1 times

⊟ **Reikidude00** 1 year, 6 months ago

Tested on GNS3, it's B

upvoted 1 times

⊟ **piojo** 1 year, 6 months ago

Selected Answer: B

Labed it, answer is B

upvoted 1 times

⊟ **larn** 1 year, 7 months ago

Selected Answer: B

In the output you can see that R1 Neig session to R2 is active, but the R2 IP is rejection the connection, therefore you need to activate the neighbor connection on router 2

upvoted 1 times

⊟ **xziomal9** 1 year, 7 months ago

Selected Answer: B

The correct answer is: B

upvoted 2 times

Question #62 Topic 1

Refer to the exhibit.

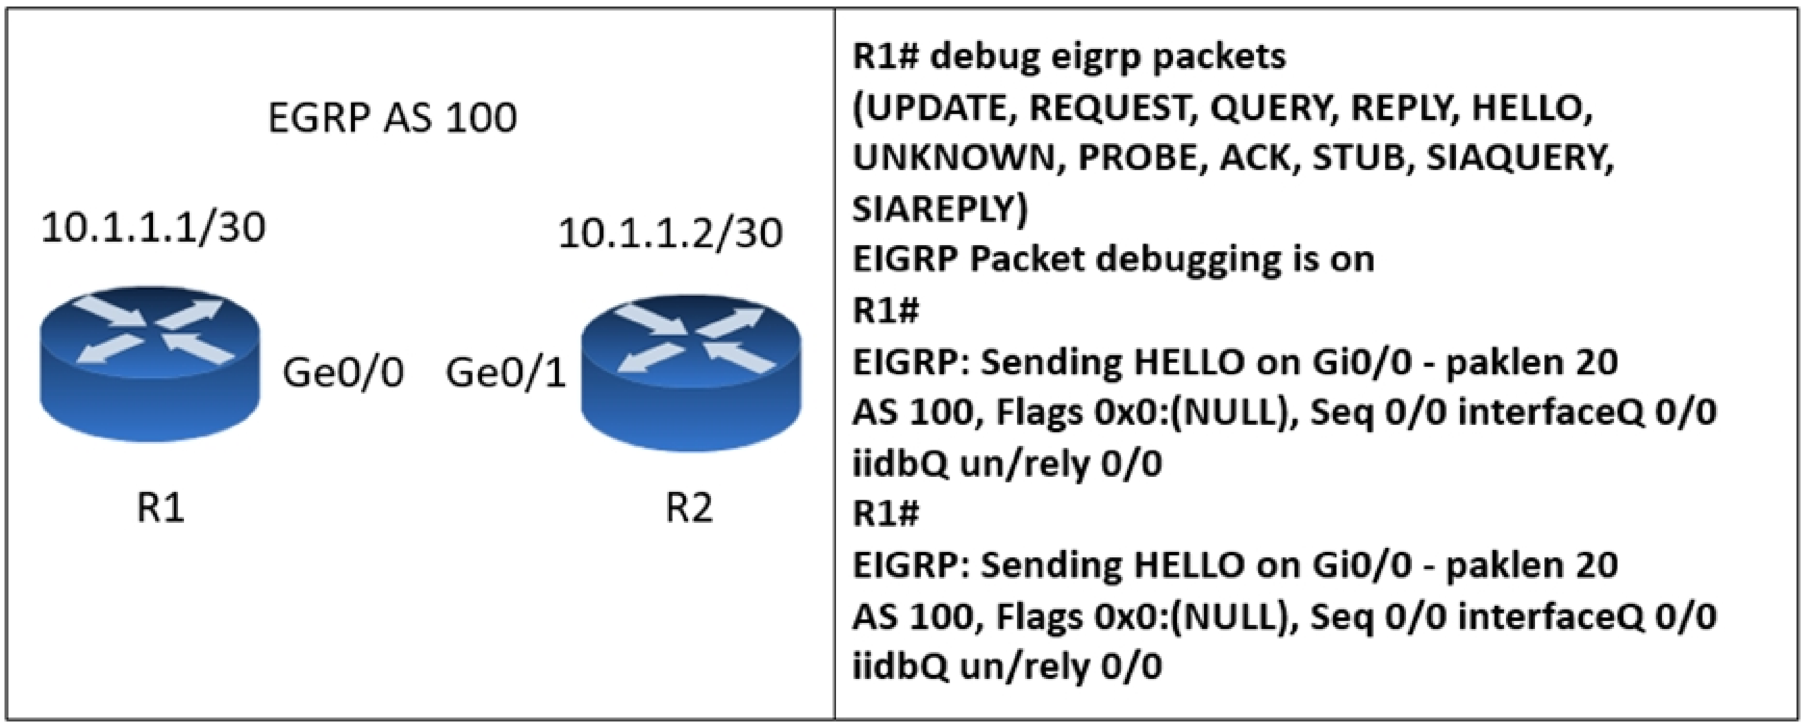

```
R1# debug eigrp packets
(UPDATE, REQUEST, QUERY, REPLY, HELLO,
UNKNOWN, PROBE, ACK, STUB, SIAQUERY,
SIAREPLY)
EIGRP Packet debugging is on
R1#
EIGRP: Sending HELLO on Gi0/0 - paklen 20
AS 100, Flags 0x0:(NULL), Seq 0/0 interfaceQ 0/0
iidbQ un/rely 0/0
R1#
EIGRP: Sending HELLO on Gi0/0 - paklen 20
AS 100, Flags 0x0:(NULL), Seq 0/0 interfaceQ 0/0
iidbQ un/rely 0/0
```

Which action resolves the adjacency issue?

A. Configure the same autonomous system numbers.

B. Match the hello interval timers.

C. Match the authentication keys.

D. Configure the same EIGRP process IDs.

Correct Answer: *A*

Reference:

https://www.ciscopress.com/articles/article.asp?p=2999383&seqNum=2

Community vote distribution

A (100%)

palihaff Highly Voted 1 year, 11 months ago

So, I tested in lab and A is correct. If auth/timers are incorrect, you will get different debug msgs. D doesn't make sense.

upvoted 5 times

SujanSikrikar Most Recent 10 months ago

Selected Answer: A

https://community.cisco.com/t5/routing/eigrp-autonomous-system-mismatch-detection/td-p/883790

upvoted 2 times

palihaff 1 year, 11 months ago

somebody, who could explain this please?

upvoted 1 times

Cyril_the_Squirl 4 months, 1 week ago

EIGRP router silently drops incoming EIGRP packets with wrong AS.

upvoted 3 times

wts 1 year, 11 months ago

Mismatched AS Numbers
When you enter the debug eigrp packets hello command, it reveals that the router does not receive the Hello packets.
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.html#anc36

upvoted 2 times

Question #63

Topic 1

Refer to the exhibit.

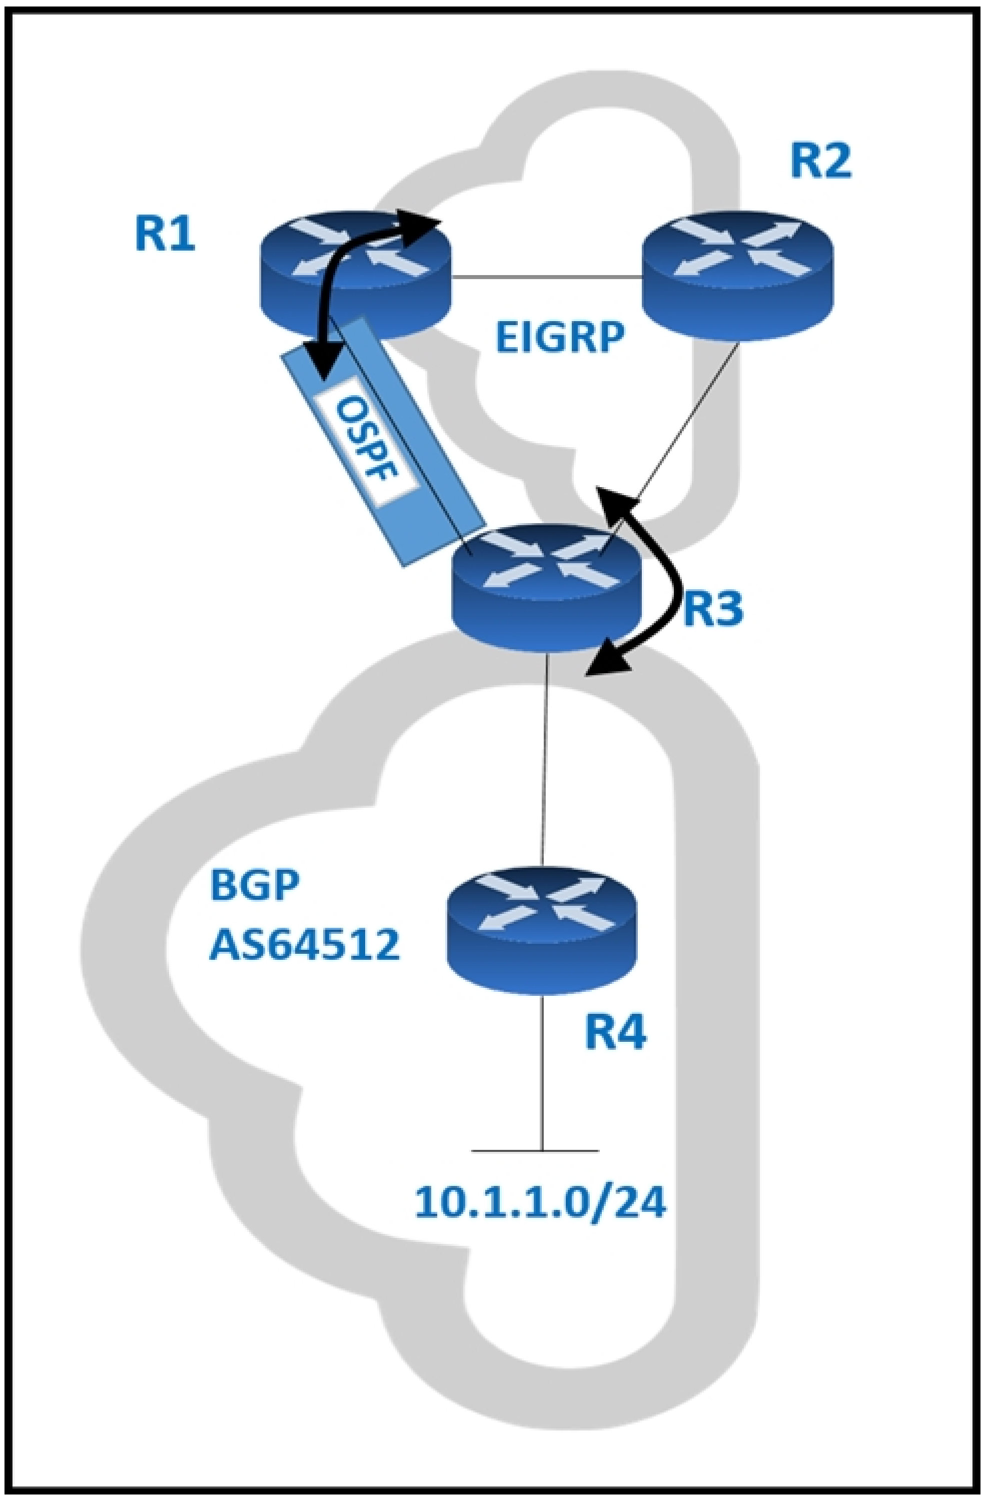

BGP and EIGRP are mutually redistributed on R3, and EIGRP and OSPF are mutually redistributed on R1. Users report packet loss and interruption of service to applications hosted on the 10.1.1.0/24 prefix. An engineer tested the link from R3 to R4 with no packet loss present but has noticed frequent routing changes on

R3 when running the debug ip route command.

Which action stabilizes the service?

A. Reduce frequent OSPF SPF calculations on R3 that cause a high CPU and packet loss on traffic traversing R3.

B. Tag the 10.1.1.0/24 prefix and deny the prefix from being redistributed into OSPF on R1.

C. Place an OSPF distribute-list outbound on R3 to block the 10.1.1.0/24 prefix from being advertised back to R3.

D. Repeat the test from R4 using ICMP ping on the local 10.1.1.0/24 prefix, and fix any Layer 2 errors on the host or switch side of the subnet.

Correct Answer: *B*

Community vote distribution

B (100%)

ciscomicha Highly Voted 1 year, 11 months ago

Given answer is correct. B is the only one with makes sence.
A = this is not an action
C = Outbound to fix advertisments back into R3? Inbound would be good but this doesn't fix the problem
D = No L2 issue

The Issue is that the AD from IBGP is 200. Highter than the AD from OSPF (110) or external EIGRP (170)

upvoted 9 times

Ll123123 Most Recent 1 month, 3 weeks ago

Selected Answer: B

Choose B. Because C apply on R3 out which is not correct

upvoted 1 times

⊟ **xziomal9** 1 year, 7 months ago

Selected Answer: B

The correct answer is: B

upvoted 1 times

Question #64 Topic 1

Refer to the exhibit. An engineer has configured policy-based routing and applied the configuration to the correct interface. How is the configuration applied to the traffic that matches the access list?

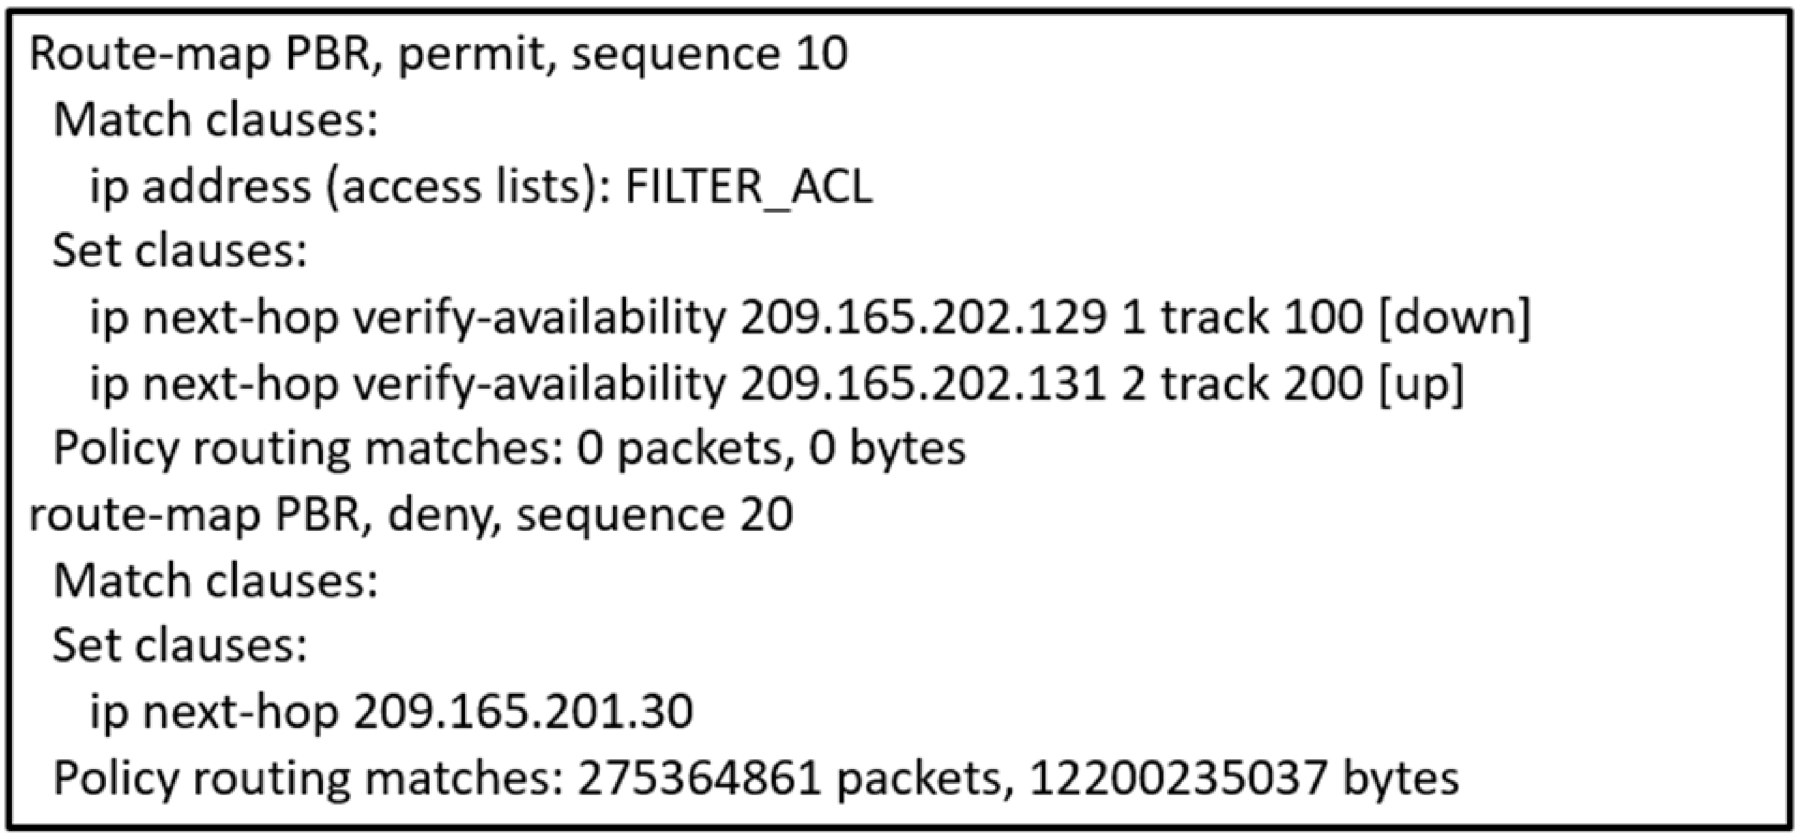

```
Route-map PBR, permit, sequence 10
  Match clauses:
    ip address (access lists): FILTER_ACL
  Set clauses:
    ip next-hop verify-availability 209.165.202.129 1 track 100 [down]
    ip next-hop verify-availability 209.165.202.131 2 track 200 [up]
  Policy routing matches: 0 packets, 0 bytes
route-map PBR, deny, sequence 20
  Match clauses:
  Set clauses:
    ip next-hop 209.165.201.30
  Policy routing matches: 275364861 packets, 12200235037 bytes
```

A. It is forwarded using the routing table lookup.

B. It is sent to 209.165.202.129.

C. It is dropped.

D. It is sent to 209.165.202.131.

Correct Answer: *D*
The first next hop IP is down, so the second one will be used.

Community vote distribution

D (100%)

JOKERR Highly Voted 1 year, 6 months ago
It's tempting to select C because of the policy routing matches. But the question explicitly states: How is the configuration applied to the traffic that matches the access list?

So the traffic matching the ACL will choose the second next-hop.

Again, trick questions designed to mess up your mind and make you fail the exam to generate revenue for Cisco.
upvoted 11 times

[Removed] Most Recent 4 months ago
Selected Answer: D
As Jokker stated. This is specific to the ACL defined, Sequence 20 does not have an ACL to match to.
Even though there are no MATCH hits on the Sequence 10, it is the only one that has an ACL.
upvoted 1 times

juliop 11 months, 2 weeks ago
The correct Anwer is A, beacause, we don´t see any maches in PBR 10 and The PBR 20 Deny statement dont mach with any ánd send the traffic for Route table.
upvoted 1 times

xziomal9 1 year, 7 months ago
Selected Answer: D
The correct answer is: D
upvoted 1 times

cyrus777 1 year, 8 months ago
Selected Answer: D
seems to be the best
upvoted 1 times

YaPet 1 year, 10 months ago

Selected Answer: D

D is correct. Here is no question how traffic will be routed. The question about how it will be routed if it is matched the ACL

upvoted 1 times

Hack4 1 year, 10 months ago

D is correct

upvoted 1 times

thinqtanklearningDOTcom 1 year, 10 months ago

We are not seeing any matched to sequence 10. 0 packets and 0 bytes, so it is matching to the deny sequence 20. Because it is a deny statement, teh set condition isn't applied and PBR is not used. Therefore we do a standard routing lookup.

upvoted 2 times

Carl1999 1 year, 10 months ago

Selected Answer: D

I think It's a Reliable PBR with IP SLA.

upvoted 2 times

ciscomicha 1 year, 11 months ago

The given answer is correct. "... matches the access list.
ACL is matched. route-map statement is permit and the second track is up. Its D

upvoted 1 times

[Removed] 1 year, 11 months ago

Doesn't matter? The question is asking *How is the configuration applied to the traffic that matches the access list* so its specifically asking about the traffic matching the acl...

upvoted 1 times

geek1992 1 year, 11 months ago

Help please we don't see match in Denison is A

upvoted 1 times

geek1992 1 year, 11 months ago

Answer is A it's match deny policy so is A

upvoted 1 times

> **cyrus777** 1 year, 8 months ago
>
> look at polcy routing matches. too many packets hit the policy
>
> upvoted 1 times

Refer to the exhibit.

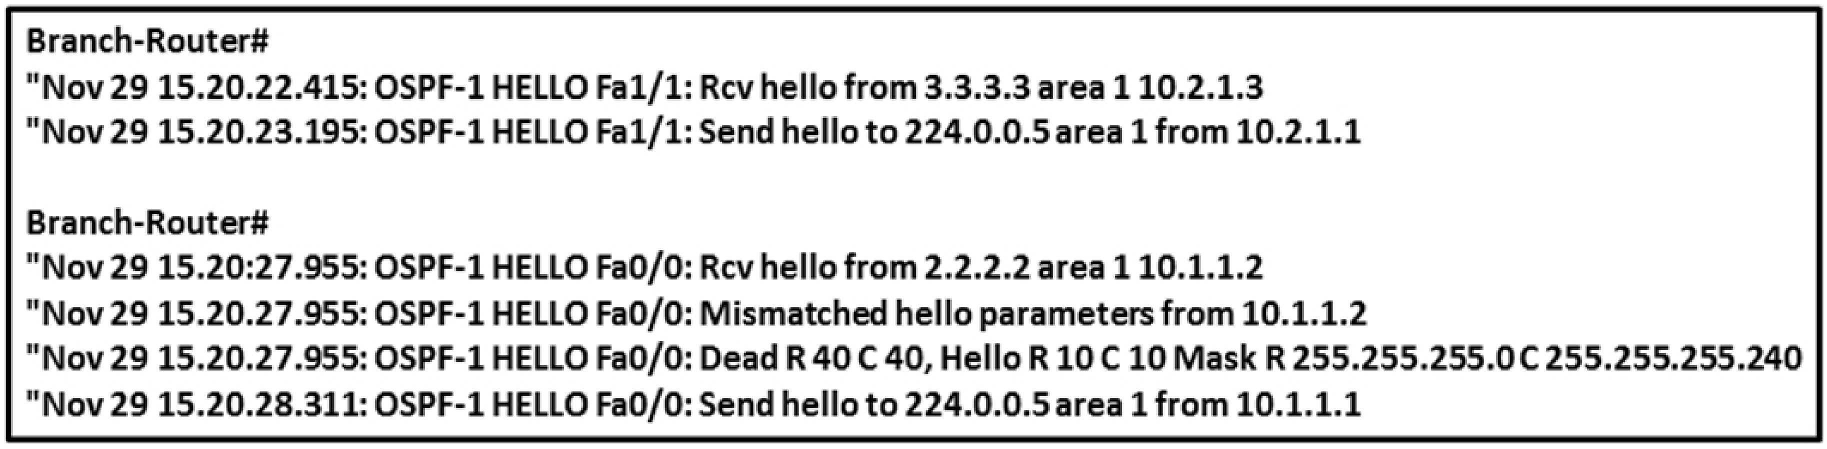

```
Branch-Router#
"Nov 29 15.20.22.415: OSPF-1 HELLO Fa1/1: Rcv hello from 3.3.3.3 area 1 10.2.1.3
"Nov 29 15.20.23.195: OSPF-1 HELLO Fa1/1: Send hello to 224.0.0.5 area 1 from 10.2.1.1

Branch-Router#
"Nov 29 15.20:27.955: OSPF-1 HELLO Fa0/0: Rcv hello from 2.2.2.2 area 1 10.1.1.2
"Nov 29 15.20.27.955: OSPF-1 HELLO Fa0/0: Mismatched hello parameters from 10.1.1.2
"Nov 29 15.20.27.955: OSPF-1 HELLO Fa0/0: Dead R 40 C 40, Hello R 10 C 10 Mask R 255.255.255.0 C 255.255.255.240
"Nov 29 15.20.28.311: OSPF-1 HELLO Fa0/0: Send hello to 224.0.0.5 area 1 from 10.1.1.1
```

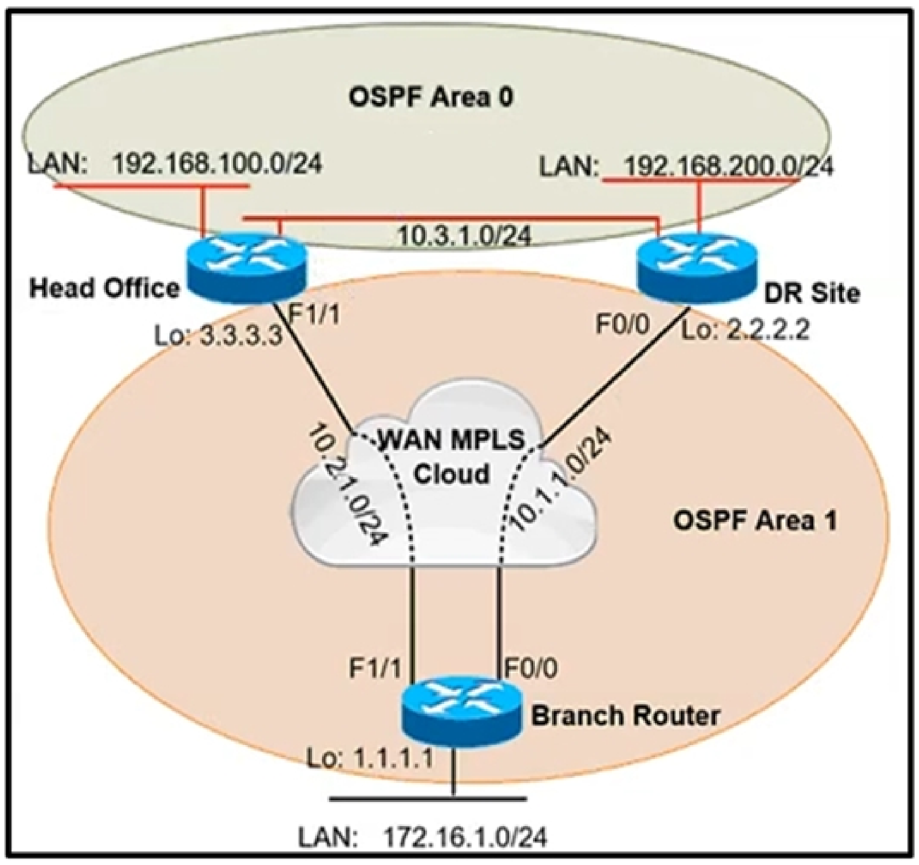

A network administrator reviews the branch router console log to troubleshoot the OSPF adjacency issue with the DR router.
Which action resolves this issue?

A. Stabilize the DR site flapping link to establish OSPF adjacency.

B. Advertise the branch WAN interface matching subnet for the DR site.

C. Configure the WAN interface for DR site in the related OSPF area.

D. Configure matching hello and dead intervals between sites.

Correct Answer: *B*
Reference:
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13699-29.html

Community vote distribution

B (92%) 8%

Edwinmolinab Highly Voted 1 year, 5 months ago

Selected Answer: B

B is more appropriate because netmask doesn't match between neighbors
R1#
OSPF: Mismatched hello parameters from 192.168.12.2
OSPF: Dead R 40 C 40, Hello R 10 C 10 Mask R 255.255.255.128 C 255.255.255.0
Now we have something to work with. R1 says it received a hello packet but we have mismatched hello parameters. The R stands for what we received and the C stands for what we have configured.

You can see that there is a mismatch in the subnet mask. R1 is configured with subnet mask 255.255.255.0 while R2 has subnet mask 255.255.255.240. OSPF will only compare the subnet mask when you are using the broadcast network type. You can also spot this error if you look at the OSPF information per interface. Broadcast is using a DR router while point-to-point doesn't use a DR router

upvoted 5 times

AinsB Most Recent 7 months, 2 weeks ago

After two routers discover each other by receiving Hellos from the other router, the routers perform the following parameter checks based on the receive Hellos:
■ Must pass the authentication process
■ Must be in the same primary subnet, including the same subnet mask
■ Must be in the same OSPF area
■ Must be of the same area type (regular, stub, not-so-stubby area [NSSA])
■ Must not have duplicate RIDs
■ OSPF Hello and Dead timers must be equal
If any of these items do not match, the two routers simply do not form a neighbor relationship.
upvoted 3 times

Dacusai 1 year, 3 months ago

Mismatched hello parameter indicates that the packet came with a wrong parameter in this case the mask addresses. timers are ok, I just lab it and get the same logs on the debug.

```
00:07:18: OSPF: Mismatched hello parameters from 10.1.1.1

00:07:18: OSPF: Dead R 40 C 40 Hello R 10 C 10 Mask R 255.255.255.0 C 255.255.255.240

00:07:28: OSPF: Rcv hello from 192.168.1.1 area 0 from GigabitEthernet0/0 10.1.1.1

00:07:28: OSPF: Mismatched hello parameters from 10.1.1.1

00:07:28: OSPF: Dead R 40 C 40 Hello R 10 C 10 Mask R 255.255.255.0 C 255.255.255.240

00:07:38: OSPF: Rcv hello from 192.168.1.1 area 0 from GigabitEthernet0/0 10.1.1.1

00:07:38: OSPF: Mismatched hello parameters from 10.1.1.1

00:07:38: OSPF: Dead R 40 C 40 Hello R 10 C 10 Mask R 255.255.255.0 C 255.255.255.240
```

upvoted 1 times

TECH3K3 1 year, 5 months ago

Selected Answer: B

The subnet mask not matching.
Hello and dead timers are fine.
upvoted 1 times

Reikidude00 1 year, 5 months ago

Selected Answer: B

Its B.
@ JingleJangus its correct.

R = Remote
C = Connected
upvoted 1 times

HungarianDish 6 months, 4 weeks ago

R = received, C = configured on the local router.
https://flylib.com/books/en/4.209.1.196/1/
upvoted 1 times

Shasha_123 1 year, 5 months ago

Selected Answer: B

It is B as it says hello parameters and not timers
upvoted 1 times

Nhan 1 year, 5 months ago

The log is clearly indicates that there is mismatch hello timer, therefore the correct answer is D
upvoted 1 times

davdtech 1 year, 6 months ago

If it was a point to point yes the subnet mask is ignored but in this case a DR has been mentioned so I would stick with the wrong subnet mask
upvoted 1 times

timtgh 1 year, 6 months ago

Selected Answer: B

B, because the subnet mask is wrong. Hello timers match.
upvoted 2 times

JOKERR 1 year, 6 months ago

But the log clearly states that Mismatched Hello Parameters from 10.1.1.2.
upvoted 1 times

JingleJangus 1 year, 6 months ago

The logs indicate a Mismatched Hello Parameter; however, WHICH parameter is mismatched?

Dead: R=40 C=40
Hello: R=10 C=10
Mask: R= -.0 C= -.240

It is the Subnet Masks that do not match. Answer is B.

upvoted 5 times

⊟ **DZhang** 1 year, 7 months ago

Selected Answer: B

dead and hello timers are same (40 and 10) but subnet is different.

upvoted 1 times

⊟ **markan** 1 year, 7 months ago

C
dead and hello timers are same (40 and 10) but Interco network different

upvoted 1 times

⊟ **xziomal9** 1 year, 7 months ago

Selected Answer: D

The correct answer is: D

upvoted 1 times

Question #66

Topic 1

Refer to the exhibit.

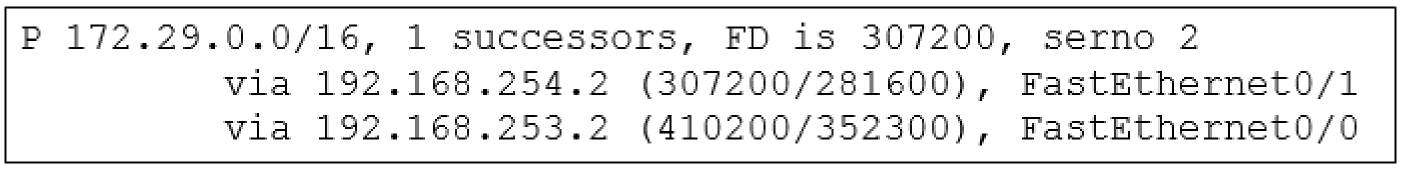

```
P 172.29.0.0/16, 1 successors, FD is 307200, serno 2
         via 192.168.254.2 (307200/281600), FastEthernet0/1
         via 192.168.253.2 (410200/352300), FastEthernet0/0
```

When the FastEthernet0/1 goes down, the route to 172.29.0.0/16 via 192.168.253.2 is not installed in the RIB. Which action resolves the issue?

A. Configure feasible distance greater than the reported distance.

B. Configure feasible distance greater than the successor's feasible distance.

C. Configure reported distance greater than the successor's feasible distance.

D. Configure reported distance greater than the feasible distance.

Correct Answer: *A*

Reference:

https://www.practicalnetworking.net/stand-alone/eigrp-feasibility-condition/

Community vote distribution

A (68%) B (32%)

Ll123123 1 month, 3 weeks ago

Selected Answer: A

I will go with A
FD = RD + local calculated metric of best route
So FD can not be configured, unless we twist the metric of successor route. And the back up reported distance must be smaller than the FD in order to be considered as backup path.

upvoted 1 times

SnoopDD 2 months ago

Selected Answer: B

For a route to be considered a backup route, the RD received for that route
must be less than the FD calculated locally. This logic guarantees a
loop-free path. FD is 307200 , RD is 352300

upvoted 1 times

Muste 4 months ago

Selected Answer: A

It should have been like this A : Configure feasible distance greater than the reported distance + of the feasible successor

upvoted 2 times

MicMillon 5 months, 3 weeks ago

Selected Answer: A

A is correct

upvoted 1 times

sajjad_gayyem 5 months, 3 weeks ago

Selected Answer: A

Only A can make change, however the answers are not precise.

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: A

For feasibility condition:
RD of feasible successor (352300) < FD of successor (307200)
1) make RD of feasible successor smaller (no such answer)
or
2) make FD of successor greater = answer "A"

upvoted 3 times

Malasxd 7 months, 1 week ago

Selected Answer: A

"A" makes more sense. If the RD in the answer is the feasible sucessor RD it is definily right.

Feasible distance = the metric of the best route (sucessor route), so B is saying for you to increase it to a value greater than itself. If you increase it to a value greater than feasible sucessor RD it would work but "A" is saying for you to do this, so it fit more. In B you can just increase the feasible distance a little bit but not enough to be greater than RD. I don't know if I was clear, my english is not that good hahaha

upvoted 2 times

6dd4aa0 8 months, 3 weeks ago

Selected Answer: B

In order to pass the feasibility condition, the feasible successor reported distance (352300) must be less than the feasibility distance (307200) in order to allow it as a backup route. In this case, it is not so.

There are two ways in doing so:
1. Lower down the feasible successor reported distance below 307200.
2. Increase the feasibility distance above the feasible successor reported distance (352300)

So, in answer B, it states to increase the feasibility distance to above the feasible successor FD (410200). As a result, it is above feasible successor reported distance (352300). This matches what I have explained in the second option.

upvoted 3 times

JoeyT 7 months ago

analyzation is correct, conclusion is wrong. bassically, you have to make 307200 bigger than 252300 or make 352300 smaller than 307200. In answers, no choice to make smaller, so you make 307200 FD bigger than 352300 RD, which is A, no doubt.

upvoted 1 times

JoeyT 7 months ago

typo, 352300. .... the other two numbers are NOT related.

upvoted 1 times

Dacusai 8 months ago

That number is already greater than the reported distance, so no make sense, answer A is more accurate making the reported distance lower than the FD.
P 10.4.4.0/24, 1 successors, FD is 3328
via 10.13.1.3 (3328/3072), GigabitEthernet0/1
via 10.14.1.4 (5376/2816), GigabitEthernet0/2
Path Metric Reported Distance
Feasible Distance
Feasible Successor
Passes Feasibility Condition
2816<3328

upvoted 3 times

davdtech 1 year, 6 months ago

Oh common, are we doing a cisco exam or a grammar exam? cisco shame on you..
So if it's answer B then it should say 'configure the FD to be greater than the feasible Suc Distance of the successor route.

upvoted 3 times

JOKERR 1 year, 6 months ago

Selected Answer: B

The answers are tricky. I am going with B because:

It says configure Feasible distance greater than successor's feasible distance. So in this case the make FD(307200) > 410200, which is greater than 352300 which would pass Feasibility condition would make the route install in the Routing table.

upvoted 2 times

sajjad_gayyem 5 months, 3 weeks ago

By doing what you say, the successor link will become the preferred route.so i go with A.

upvoted 1 times

Koume 11 months, 1 week ago

Remember that the feability condition is "The router Reported distance should be less than the successor feasible distance. the only Feasible distance that you can change is the succesor FD to make the second route meet the criteria.

upvoted 2 times

timtgh 1 year, 6 months ago

I have a theory. I believe there are two typos.
First, the question should say "what is wrong?" and not "what action will fix it."
Second, remove the word "configure" from the answers.
Then one answer makes sense: C. The problem is the reported distance is greater than the successor's FD.

upvoted 2 times

jester_2020 1 year, 7 months ago

The question is confusing and kinda gramatically incorrect. According to Feasible Condition, the RD of Feasible Successor must be lower or less than the FD of successor. Based on the question, it's not clear which metric to change, the sucessor or the feasible successor?

upvoted 1 times

Hack4 1 year, 10 months ago

A is correct

upvoted 1 times

Carl1999 1 year, 10 months ago

Selected Answer: A

A is correct:

upvoted 3 times

Carl1999 1 year, 10 months ago

This is not good question.

upvoted 2 times

krn007 1 year, 11 months ago

Suggested Answer A is correct:
Feasibility condition for EIGRP is Reported Distance (RD) of the possible feasible successor should be less than Feasibility Distance (FD) of the successor Route (i.e RD<FD). in the exhibit (RD=352300 > FD=307200) and hence the route to 172.29.0.0/16 via 192.168.253.2 is NOT installed in the RIB; however it will be available in eigrp topology table (accessible with show ip eigrp topology all-links).
One the Feasibility condition is passed by configuring FD>RD and the route will become displayed in routing table.
please let me know further comments if any.

upvoted 4 times

Networkingguy 1 year, 11 months ago

I am not sure here, could it not be B? Configure feasible distance greater than the successor's feasible distance.

upvoted 1 times

Networkingguy 1 year, 11 months ago

So if the AD for the non successor route is less than the FD of the successor, the route is a feasible successor. Which in this case it is, perfect 352300 is the AD (Reported/Advertised Distance) and 410200 is its FD (Feasible Distance). Answer A is already true... A. Configure feasible distance greater than the reported distance.

Options B, C and D are all also true. Strange question?

upvoted 1 times

Networkingguy 1 year, 11 months ago

Do we manually need to change the FD when a link changes? Would the router hold onto the 307200 and thus would freak out about the feasible successor's new path. It must..

upvoted 1 times

[Removed] 1 year, 11 months ago

Change the variance.. It will then multiply the FD by whatever you choose between 1-4.

upvoted 2 times

[Removed] 1 year, 11 months ago

Huh? The reported distance is 352300 and the feasible distance is 307200. In order for a route to be considered a feasible successor it has to meet the feasibility condition of FD > RD and in this case the RD is higher. So A isn't true and needs to be done for the route to be installed into the RIB...

upvoted 2 times

Refer to the exhibit.

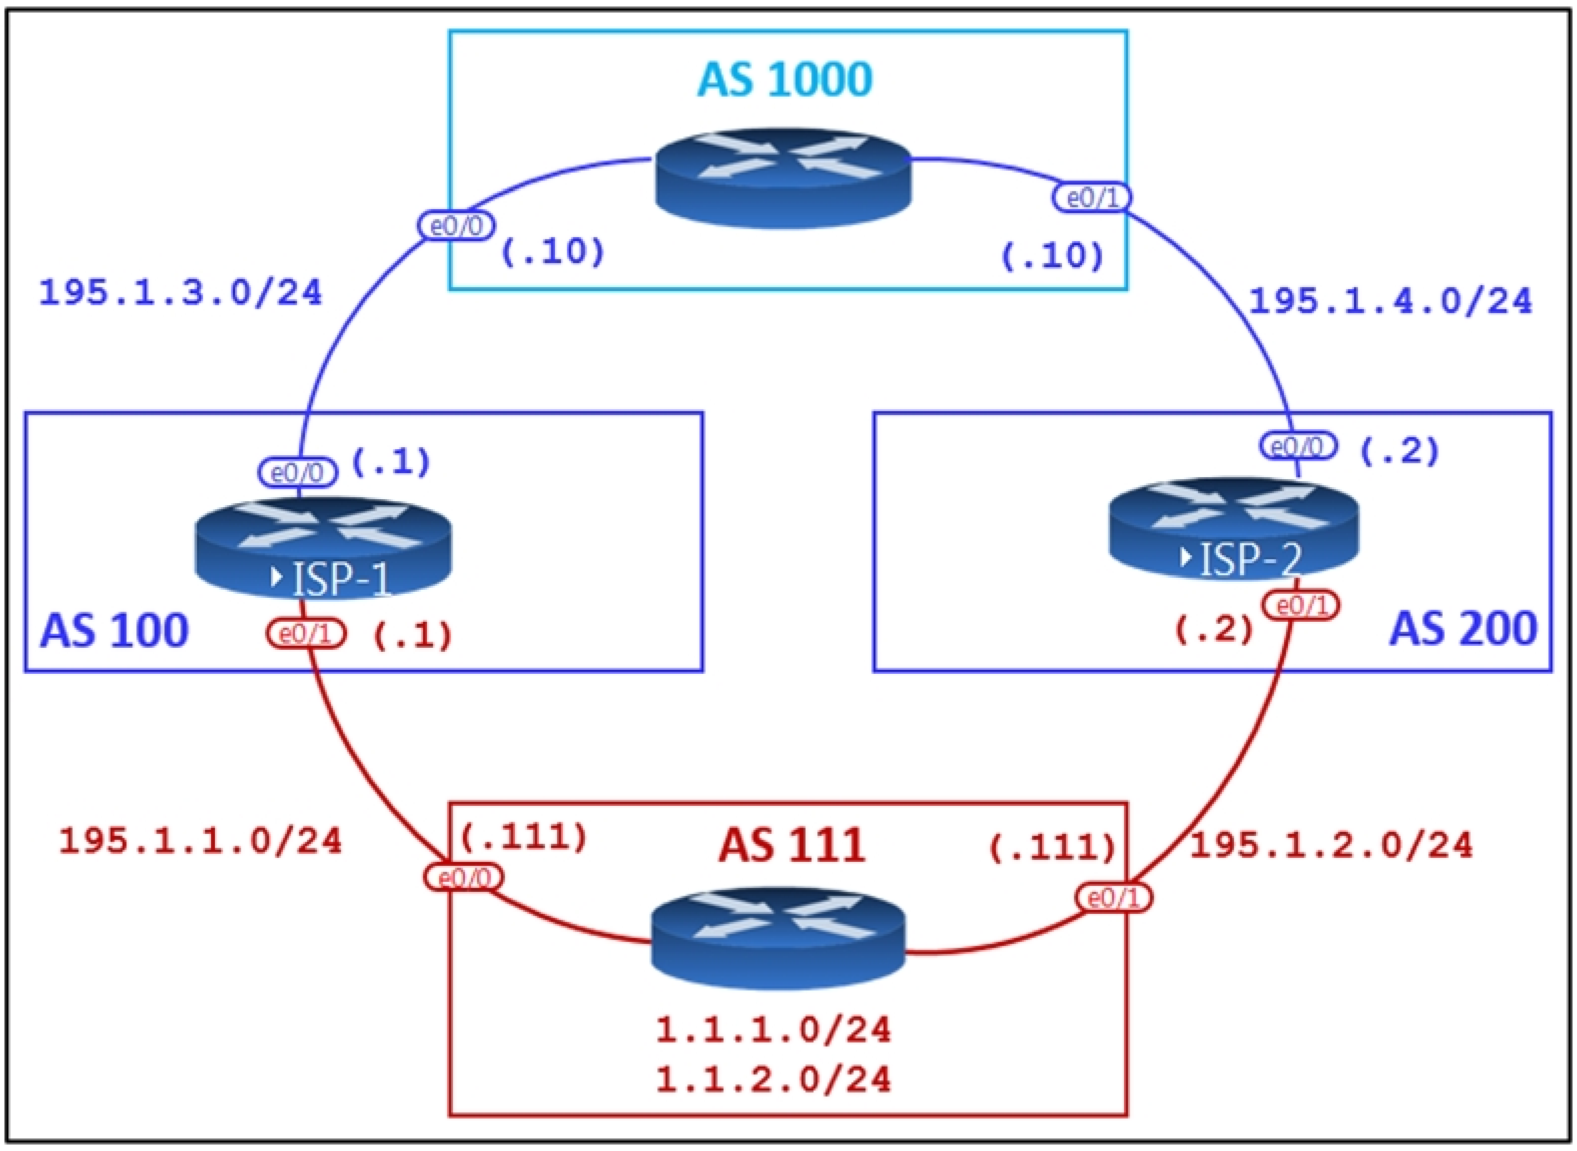

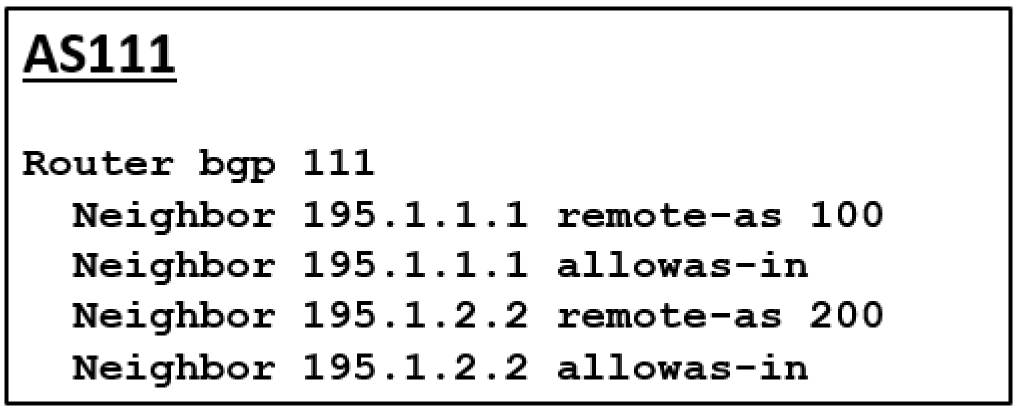

AS111

```
Router bgp 111
  Neighbor 195.1.1.1 remote-as 100
  Neighbor 195.1.1.1 allowas-in
  Neighbor 195.1.2.2 remote-as 200
  Neighbor 195.1.2.2 allowas-in
```

AS111 is receiving its own routes from AS200 causing a loop in the network.
Which configuration provides loop prevention?

A. router bgp 111 neighbor 195.1.1.1 as-override no neighbor 195.1.2.2 allowas-in

B. router bgp 111 no neighbor 195.1.1.1 allowas-in no neighbor 195.1.2.2 allowas-in

C. router bgp 111 neighbor 195.1.2.2 as-override no neighbor 195.1.1.1 allowas-in

D. router bgp 111 neighbor 195.1.1.1 as-override neighbor 195.1.2.2 as-override

Correct Answer: *B*

Reference:

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/112236-allowas-in-bgp-config-example.html

Community vote distribution

B (100%)

Colmenarez 3 months, 3 weeks ago

This reminds me of an encore lab question.

upvoted 2 times

xziomal9 1 year, 7 months ago

Selected Answer: B

The correct answer is: B

router bgp 111

no neighbor 195.1.1.1 allowas-in
no neighbor 195.1.2.2 allowas-in

upvoted 2 times

⊟ **Carl1999** 1 year, 10 months ago

B is correct.
In this case, Allow as-in is not needed.
Because it disables control of AS_PATH.

upvoted 2 times

Question #68 Topic 1

Refer to the exhibit.

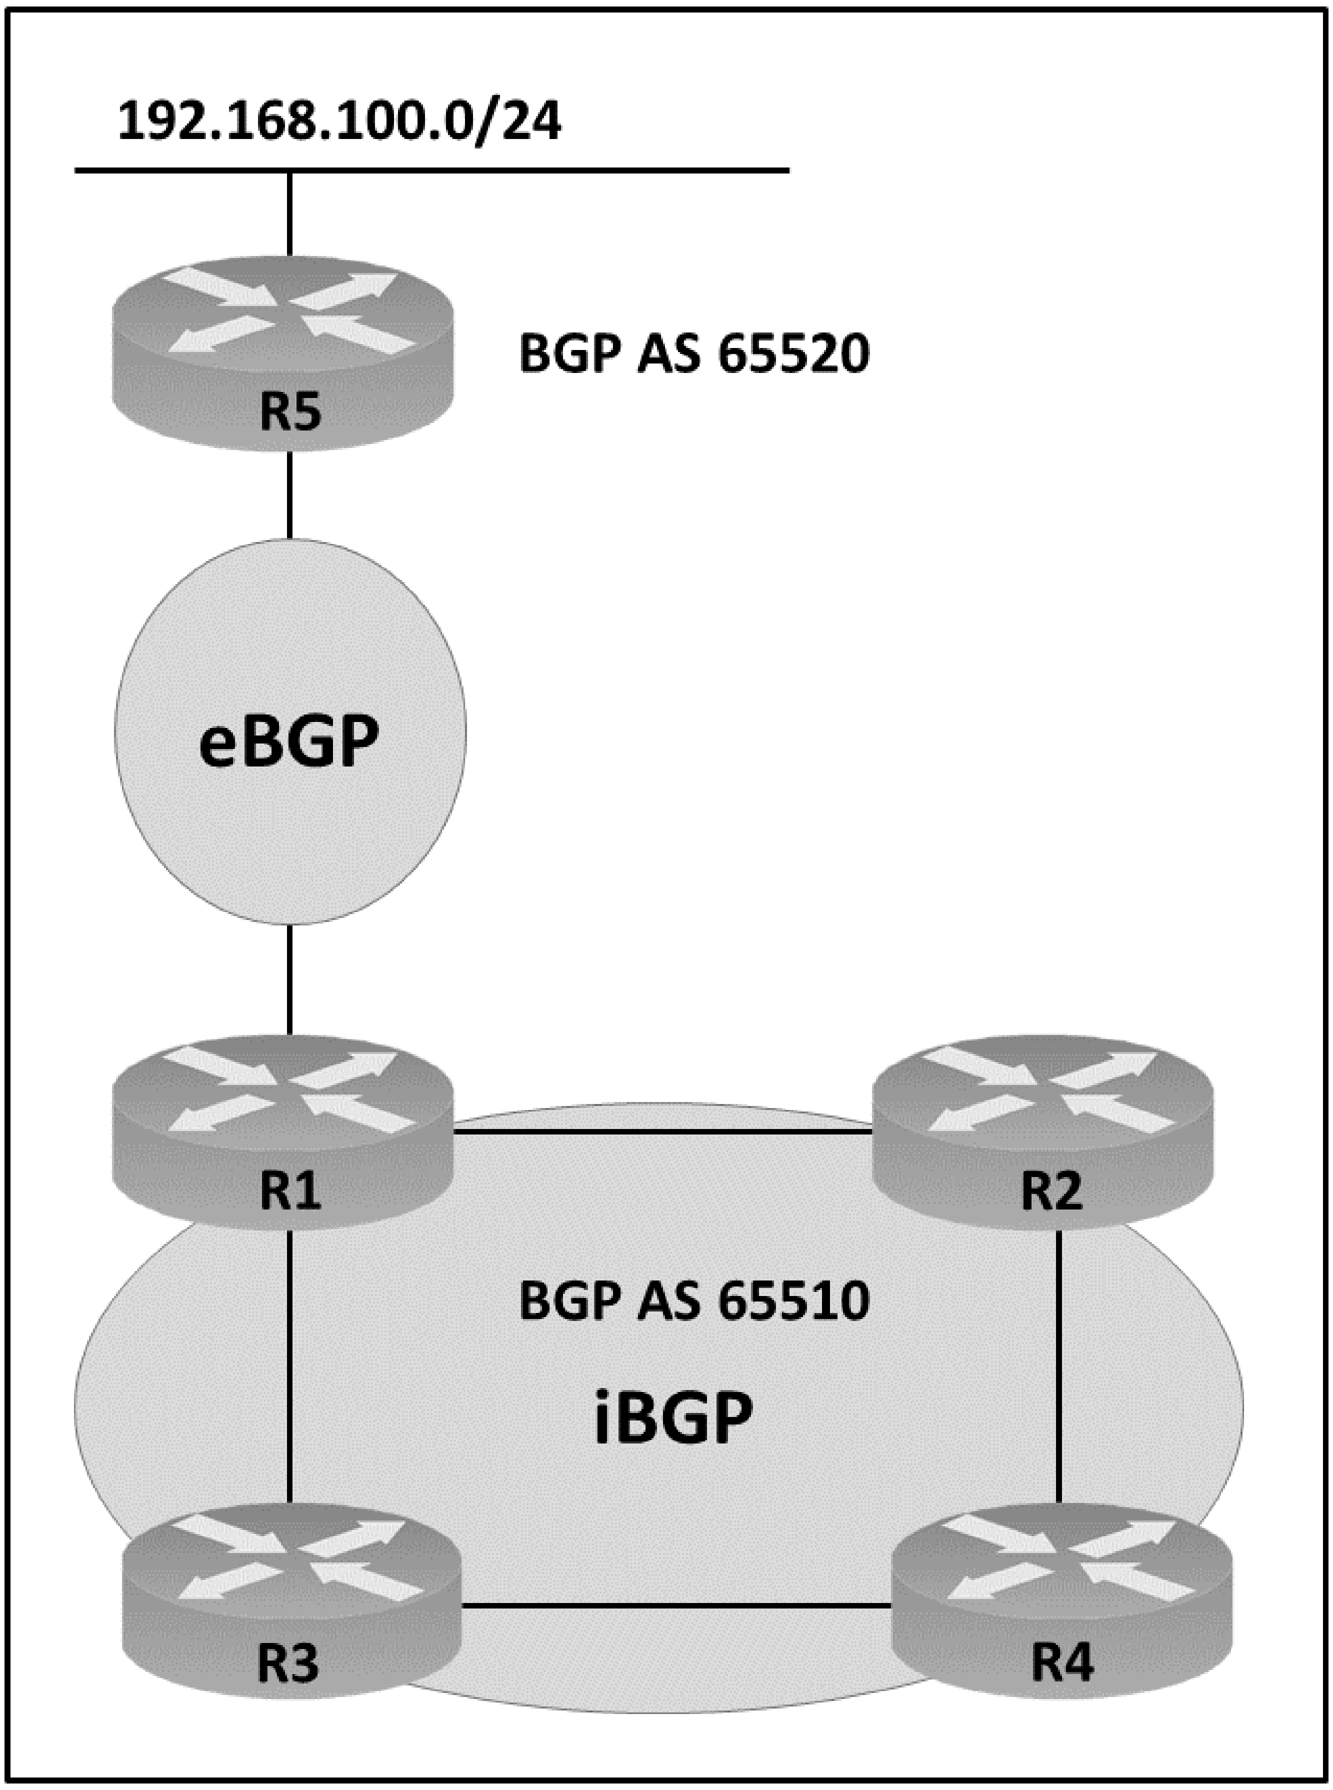

AS65510 iBGP is configured for directly connected neighbors. R4 cannot ping or traceroute network 192.168.100.0/24.
Which action resolves this issue?

A. Configure R1 as a route reflector server and configure R2 and R3 as route reflector clients.

B. Configure R4 as a route reflector server and configure R2 and R3 as route reflector clients.

C. Configure R4 as a route reflector server and configure R1 as a route reflector client.

D. Configure R1 as a route reflector server and configure R4 as a route reflector client.

Correct Answer: *D*

Community vote distribution

B (39%) D (39%) C (23%)

⊟ **timtgh** Highly Voted 1 year, 6 months ago

All answers are wrong. R1 and R4 do not see each other's routes because they are two hops apart. They need a route reflector between them, either R2 or R3.

upvoted 14 times

⊟ **Pietjeplukgeluk** 1 month, 1 week ago

Indeed the quality of the question is low. To have any solution work, you need to add neighbor config on top of the specified "directly connected". Spend you time well and understand why all are wrong.
upvoted 1 times

larn Highly Voted 1 year, 7 months ago

Selected Answer: B

Dont think it possible to be D as the question clearly states there is ONLY BGP relationship between directly connected devices, D is only possible if you stand a BGP relationship between R1 & R4, then you have a full mesh and dont need route reflector at all.

If R2 & R3 are RR clients and R4 The RR. R1 advertises to R2 & R3 they then pas the route to the reflector.
upvoted 6 times

glbngl91 10 months, 1 week ago

Ehm... wrong, if you have connectivity between loopback via an IGP (or static routing), you can configure peering between two routers that are not directly connected... in my opinion D is not correct, but only because R1 is a border router and it's better not to configure it as a RR
upvoted 1 times

tinoe Most Recent 1 day, 15 hours ago

R4 can only be a route reflector client of either R2 or R3 for the ping to work. This means for this scenario question there is no correct answer, all the answers are incorrect.
upvoted 1 times

DeWalt95 2 weeks, 1 day ago

D is the only one that works but even then only if there is a direct peering between the routers (and the diagram implies there isnt).
upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

The anwser correct is D, because, router 1 y router 4 have to be RR
upvoted 1 times

SolidSnake74 5 months, 3 weeks ago

Answer is A
R4 can't reach 192.168.100.0/24
It means that R4 has most likely no routes received from any BGP neighbors.
So, by setting R1 as RR with R2 & R3 as clients, they will get the routes from remote AS via R1.
They can, then, just share it with their local neighbor which is R4.

Route reflector rule 3 says : If an RR receives a route from an eBGP peer, it advertises the route to RR clients and non-RR clients
upvoted 1 times

yasiramith 6 months ago

Selected Answer: D

D is correct
upvoted 1 times

shiro990127 6 months, 3 weeks ago

Selected Answer: C

according to Cisco, R1 has to be RR client.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/212881-border-gateway-protocol-bgp-optimal-ro.html
-The border routers must be RR clients of the RR.
upvoted 1 times

HungarianDish 6 months, 4 weeks ago

I came to the same conclusion as others. Some information is missing from this question, or the question is different in the real exam.
I tested all "A", "B", "C", "D" in CML. The design of "A" and "B" does not work. "C" and "D" worked, of course, but not because of the RR configuration.

We need to avoid setting the border router (R1) as RR. So, we can exclude "A" and "D".
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/212881-border-gateway-protocol-bgp-optimal-ro.html
upvoted 2 times

HungarianDish 6 months, 4 weeks ago

Normally, R2 OR R3 are the best candidates for becoming RR (the desired connection works without any further configurations), but such option is not given.
If setting up R1 or R4 as the RR, we need to add IGP or static routes, so R1 and R4 can become neighbors first.
In this case, we build a full mesh and we do not need any RR.

The question says: "iBGP is configured for directly connected neighbors".
-> It does not mention, but does not exclude either an underlay routing being configured for R1 and R4 neighborship. So, it might be a full mesh.
upvoted 1 times

HungarianDish 6 months, 4 weeks ago

+R1 needs next-hop-self for sure.

upvoted 1 times

AinsB 7 months, 1 week ago

Selected Answer: D

IBGP will not readvertise a BGP learnt route to another IBGP neighbor, that is why we need a RR so that the route can get to R4. If R1 is the RR it can be advertised to R4 because a TCP connection will be created. Based on this scenario R2 & R3 could be RR and accomplish the same thing but they are not a part of the answer , it would be a much better design though

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Solution "C" has the same result, but it has the advantage that it does not use the border router R1 as an RR. The thing is that: If we make R1 and R4 become neighbors, and configure next-hop-self on R1, then it works without RR. Only R2 or R3 work well as RR.

upvoted 1 times

Dacusai 8 months ago

I don't see a correct answer here, R2 and R3 will received routes from eBGP on R1 because is external so for R4 get those routes and R1 get routes from R4 you need to make them RRC on R2 and R3 with the command neighbor (ip) route-reflector-client for both R1 and R4. In that way R2 and R3 will send the routes to R4 and R1.

upvoted 2 times

forccnp 9 months ago

Selected Answer: B

i vote for B

upvoted 1 times

glbngl91 10 months, 1 week ago

Selected Answer: C

Guys, it's C in my opinion. I tried option A and B on GNS3 and both do not work, and this is obvious since R2 and R3 are not aware of being RR-client, so split horizon is enabled and they won't announce to R4 (answer A) or R1 (answer B) a network learned by another ibgp peer... Answer C and D both work, but the best practise is not to configure the RR on a border router, so for me it's C

upvoted 1 times

toto89 11 months, 2 weeks ago

Selected Answer: C

For me C and D works, but C is better since its better not to place the RR on a boundary router.

upvoted 2 times

christiannomarcenes 1 year, 3 months ago

I think it is D. R2 and R3 being server is implicit (you have to imagine this). I tested GNS all choices and D is the only that works.

upvoted 4 times

doron1122 1 year, 4 months ago

i think b

r1 = EBGp
2+3 ebgp to IBGP
4 needs to be RR to get from R2 & R3
So i choose B

upvoted 2 times

Reikidude00 1 year, 5 months ago

Selected Answer: B

B is correct because is not the best practise to make a border router a RR

upvoted 1 times

ericxw 11 months, 3 weeks ago

when u do cisco exams no answer is best practice

upvoted 5 times

Question #69 Topic 1

Users report issues with reachability between areas as soon as an engineer configured summary routes between areas in a multiple area OSPF autonomous system.
Which action resolves the issue?

A. Configure the area range command on the ASBR.

B. Configure the summary-address command on the ASBR.

C. Configure the summary-address command on the ABR.

D. Configure the area range command on the ABR.

Correct Answer: *D*

Community vote distribution

D (100%)

Nhan 1 year, 5 months ago

D is correct answer

The area range command is used only with area border routers (ABRs). It is used to consolidate or summarize routes for an area. The result is that a single summary route is advertised to other areas by the ABR. Routing information is condensed at area boundaries. External to the area, a single route is advertised for each address range.

upvoted 1 times

piojo 1 year, 6 months ago

Answer should be:

D. Configure the area range command on the ABR CORRECTLY

upvoted 1 times

timtgh 1 year, 6 months ago

All answers are wrong. A and C have invalid syntax. B is for ASBRs, not relevant here. D is the only command that makes sense, but that is the command that was already used to configure the summarization.

upvoted 1 times

xziomal9 1 year, 7 months ago

Selected Answer: D

The correct answer is: D

upvoted 1 times

Question #70 Topic 1

Refer to the exhibit.

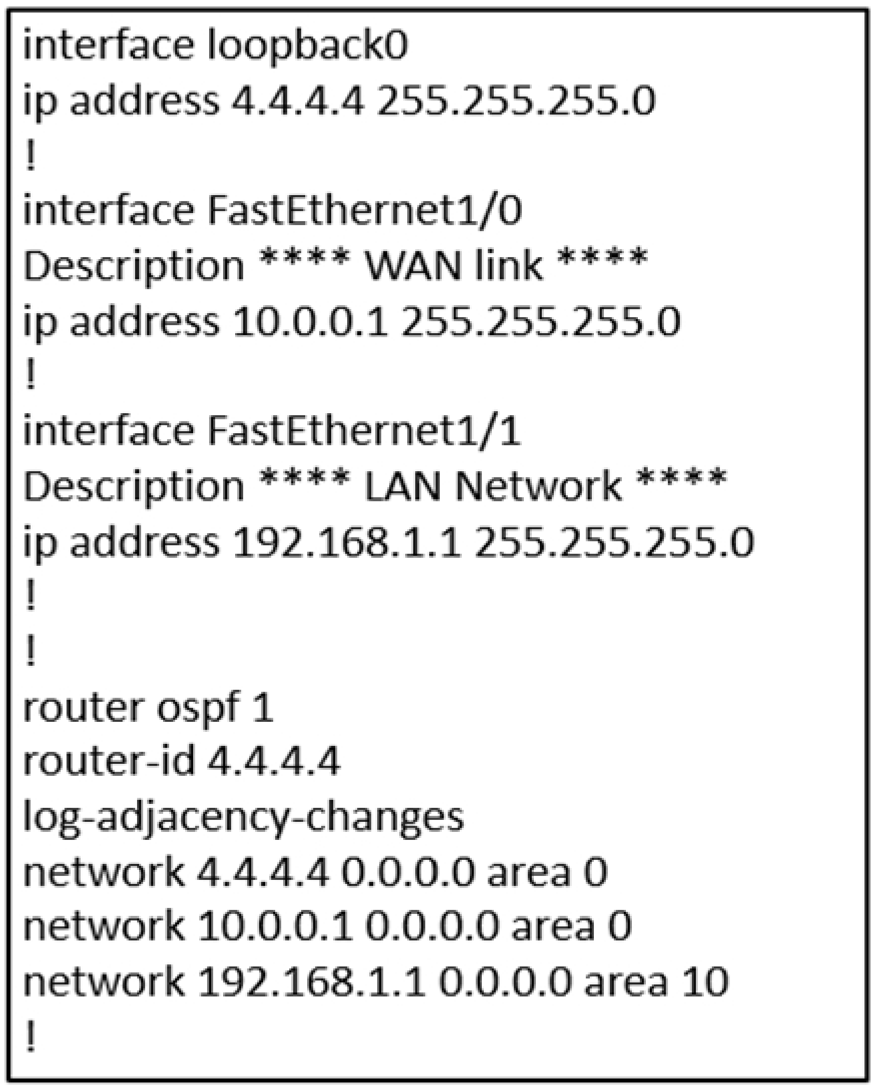

```
interface loopback0
ip address 4.4.4.4 255.255.255.0
!
interface FastEthernet1/0
Description **** WAN link ****
ip address 10.0.0.1 255.255.255.0
!
interface FastEthernet1/1
Description **** LAN Network ****
ip address 192.168.1.1 255.255.255.0
!
!
router ospf 1
router-id 4.4.4.4
log-adjacency-changes
network 4.4.4.4 0.0.0.0 area 0
network 10.0.0.1 0.0.0.0 area 0
network 192.168.1.1 0.0.0.0 area 10
!
```

Which set of commands restore reachability to loopback0?

A. interface loopback0 ip address 4.4.4.4 255.255.255.0 ip ospf network point-to-point

B. interface loopback0 ip address 4.4.4.4 255.255.255.0 ip ospf interface area 10

C. interface loopback0 ip address 4.4.4.4 255.255.255.0 ip ospf network broadcast

D. interface loopback0 ip address 4.4.4.4 255.255.255.0 ip ospf interface type network

Correct Answer: *A*

Reference:

https://networkengineering.stackexchange.com/questions/13099/why-do-we-use-ospf-point-to-point-networks-for-loopbacks

Community vote distribution

A (100%)

heeeeyajoke 1 year ago

I have tested the lab, i could still reach the loopback, OSPF advertised a /32 route for the loopback by default. But i would go for the option A for the exam

upvoted 2 times

anaisa_goncalves 1 year, 1 month ago

Here's the explanation:
https://networkengineering.stackexchange.com/questions/13099/why-do-we-use-ospf-point-to-point-networks-for-loopbacks

upvoted 1 times

Carl1999 1 year, 10 months ago

Selected Answer: A

Loopbacks are considered host routes in Open Shortest Path First (OSPF).
To make OSPF advertise the loopback subnet as the actual subnet with the loopback mask, instead of as host route /32, issue the ip ospf network point-to-point command under the loopback interface.
For more information, refer to the 9.1. Interface States section of RFC 2328.

upvoted 3 times

diogodds 1 year, 9 months ago

The question is really horrible as mentioned by @bogd, with the current interface and ospf configuration, the loopback would also be reachable.

upvoted 2 times

Pietjeplukgeluk 1 month ago

Fully correct, the loopback would be reachable however it would advertise /32 instead of a /24. So there is no need to correct any configuration, changing it would only advertise different mask, but you would not be able to reach any other interface in the actual advertised network. Stupid question.

upvoted 1 times

palihaff 1 year, 11 months ago

I understand the answ, but please, somebody explain, to me how the original config can make 4.4.4.4 unreachable. Thanks !

upvoted 3 times

OhBee 1 year, 11 months ago

I believe it is because the route is advertised as a /32 (since OSPF does that by default for loopback interfaces). In order for it to be advertised properly as a /24, the configuration shown in A should be done.

upvoted 2 times

bogd 1 year, 10 months ago

Having the route advertised as a /24 or /32 really doesn't matter - 4.4.4.4 would be reachable in both cases. The phrasing of the question is absolutely terrible...

upvoted 5 times

piojo 1 year, 6 months ago

Point is that if you configured the loopback as a /24 is because you need the whole /24 to be advertised as well (for other purposes like NATed addresses). Otherwise you would configure the loopback as /32.

upvoted 1 times

Question #71 *Topic 1*

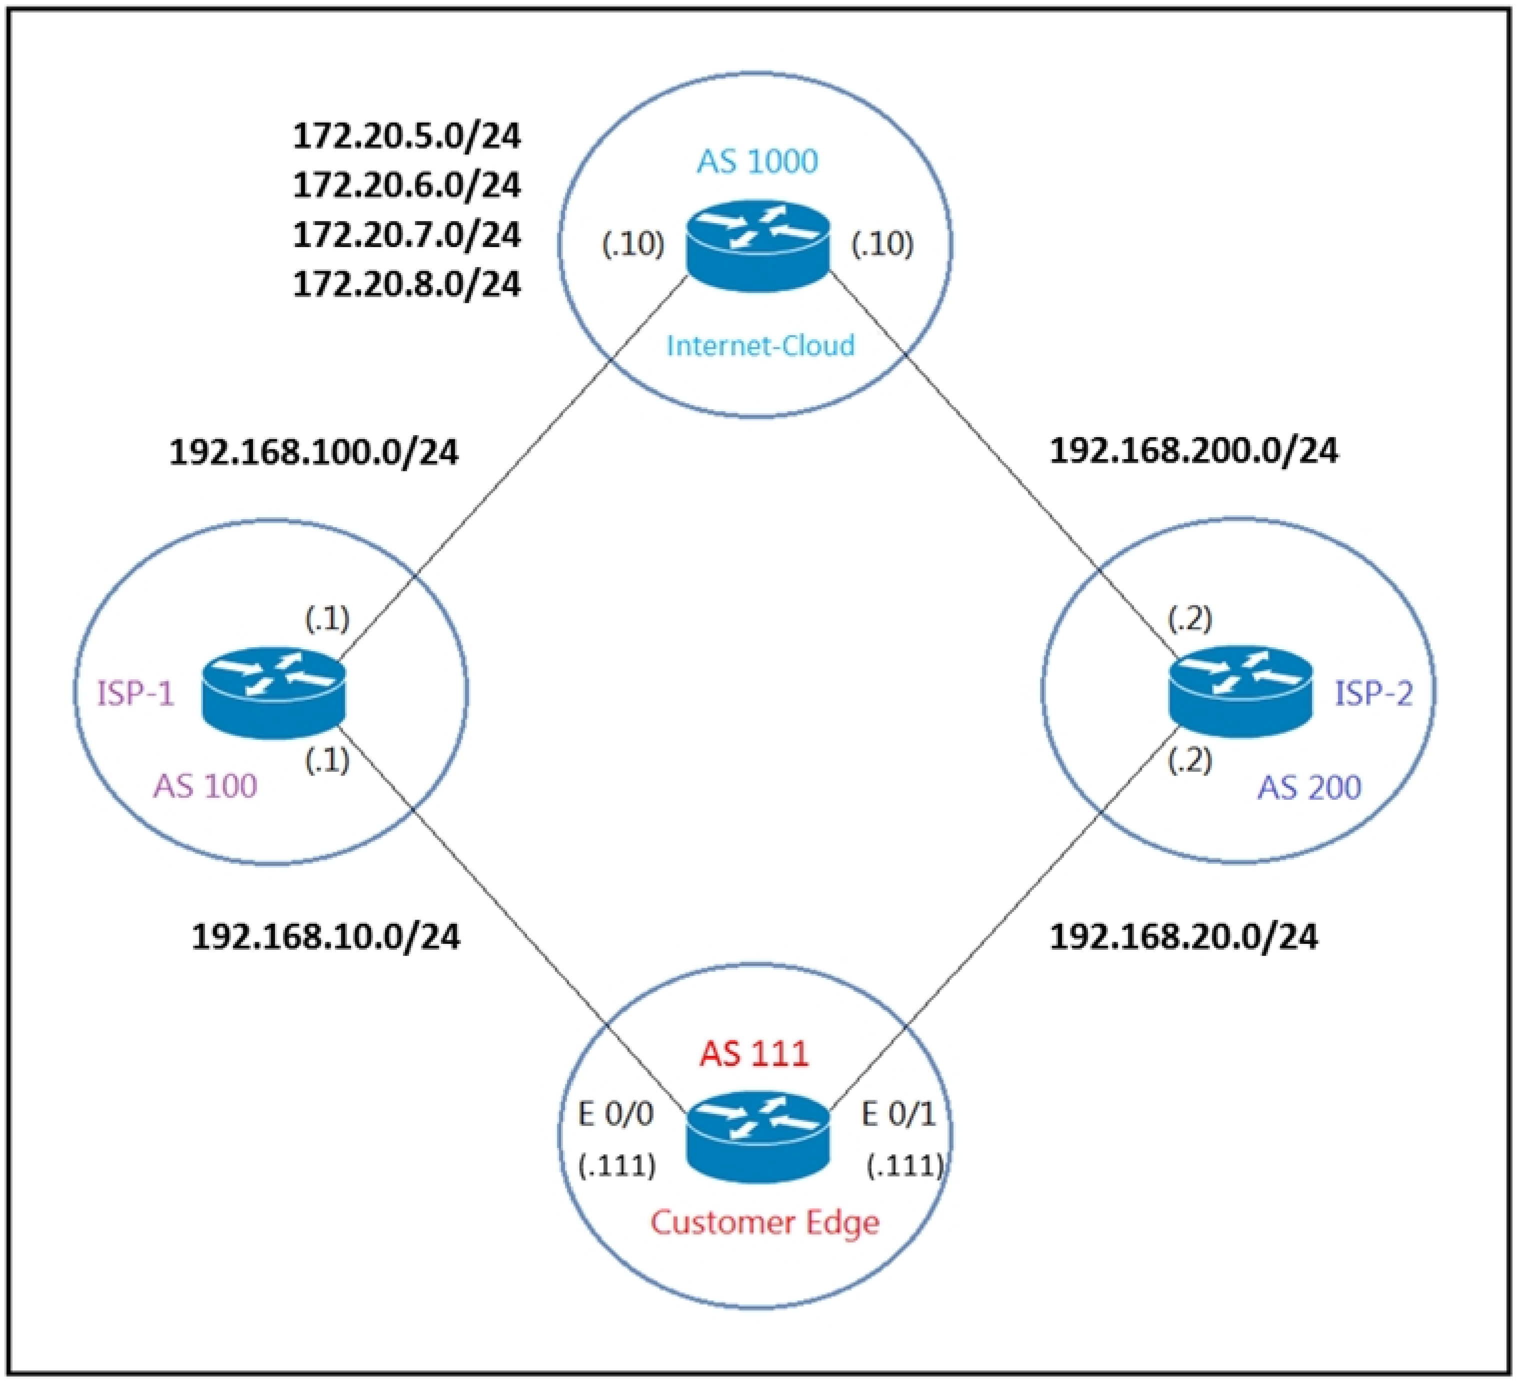

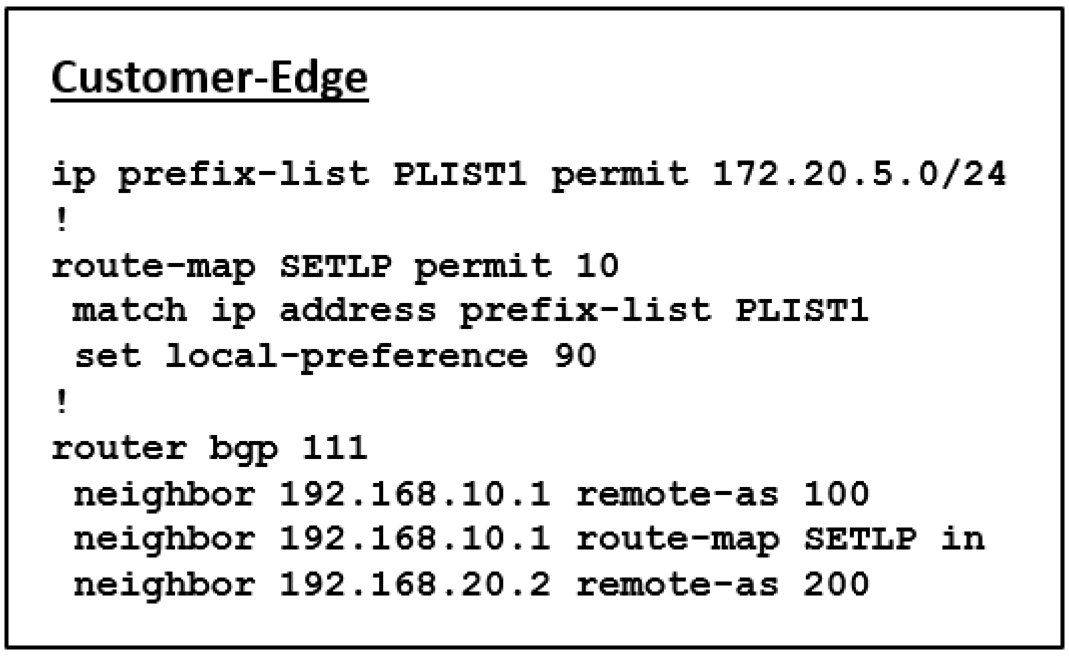

Customer-Edge

```
ip prefix-list PLIST1 permit 172.20.5.0/24
!
route-map SETLP permit 10
 match ip address prefix-list PLIST1
 set local-preference 90
!
router bgp 111
 neighbor 192.168.10.1 remote-as 100
 neighbor 192.168.10.1 route-map SETLP in
 neighbor 192.168.20.2 remote-as 200
```

AS 111 wanted to use AS 200 as the preferred path for 172.20.5.0/24 and AS 100 as the backup. After the configuration, AS 100 is not used for any other routes.

Which configuration resolves the issue?

A. route-map SETLP permit 10 match ip address prefix-list PLIST1 set local-preference 99 route-map SETLP permit 20

B. router bgp 111 no neighbor 192.168.10.1 route-map SETLP in neighbor 192.168.20.2 route-map SETLP in

C. route-map SETLP permit 10 match ip address prefix-list PLIST1 set local-preference 110 route-map SETLP permit 20

D. router bgp 111 no neighbor 192.168.10.1 route-map SETLP in neighbor 192.168.10.1 route-map SETLP out

Correct Answer: *A*

There is an implicit deny all at the end of any route-map so all other traffic that does not match 172.20.5.0/24 would be dropped. Therefore, we have to add a permit sequence at the end of the route-map to allow other traffic.

The default value of Local Preference is 100 and higher value is preferred so we have to set the local preference of AS100 lower than that of AS200.

Community vote distribution

A (91%) 9%

JingleJangus Highly Voted 1 year, 10 months ago

Selected Answer: A

A works because the default local pref is 100. Making the local pref 99 will decrement the quality of the route enough to install the route from AS200 into the RIB.

The reason traffic wasnt using AS100 for any other routes is because of the logic of the route map. Seq 10 matches on a specific range of ip's. Therefore traffic that doesnt match that range will move onto the next seq #, but in our case there is none so all other traffic hits the implicit deny at the end of the route map. Adding a seq 20 to match on all traffic (not having a match statement) will allow all other NLRI thru from AS100 to AS111.

upvoted 10 times

AinsB Most Recent 7 months, 2 weeks ago

Selected Answer: C

correct answer is C, higher local preference is preferred

upvoted 1 times

AinsB 7 months ago

I am correcting my answer to "A" for the reason that we need to route to go to AS200 so if we drop the local preference to 99 for AS100 the route will prefer AS200 "BUT" note that there is an implicit deny at the end of the prefix list/access list that we need to override so that other routes an take this path. This is the trick of these exam questions "good review"

upvoted 1 times

Noproblem22 1 year, 1 month ago

A is correct

upvoted 1 times

Question #72 *Topic 1*

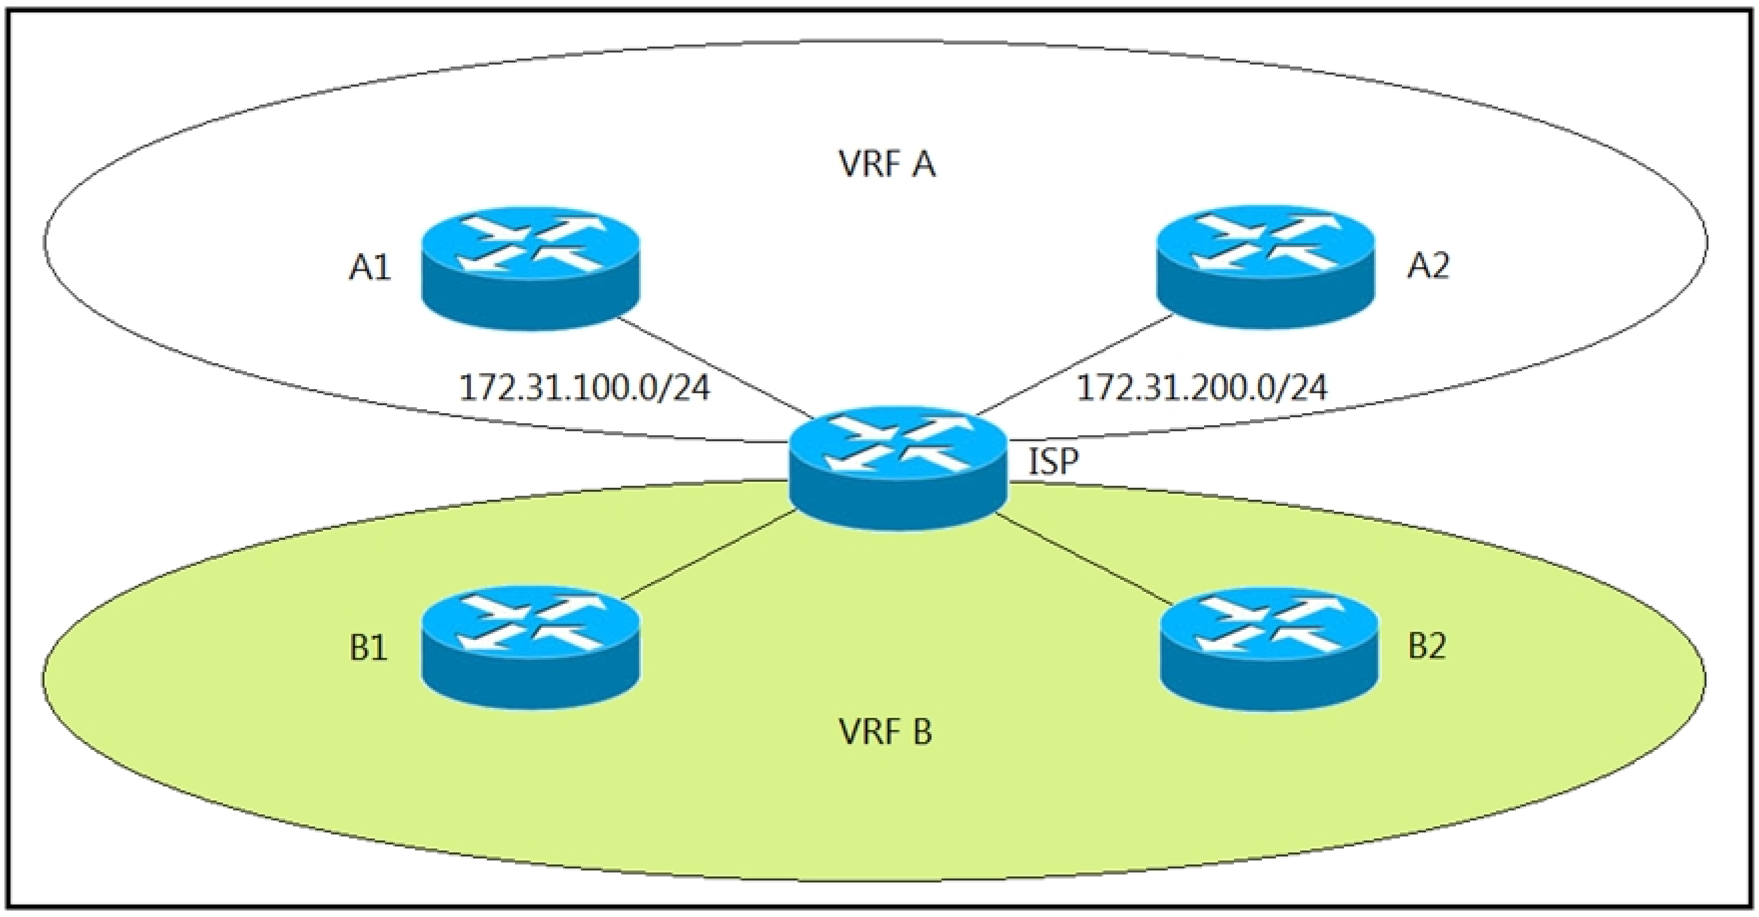

Refer to the exhibit. The ISP router is fully configured for customer A and customer B using the VRF-Lite feature.
What is the minimum configuration required for customer A to communicate between routers A1 and A2?

A. A1 interface fa0/0 description To->ISP ip add 172.31.100.1 255.255.255.0 no shut ! router ospf 100 net 172.31.100.1 0.0.0.255 area 0 A2 interface fa0/0 description To->ISP ip add 172.31.200.1 255.255.255.0 no shut ! router ospf 100 net 172.31.200.1 0.0.0.255 area 0

B. A1 interface fa0/0 description To->ISP ip vrf forwarding A ip add 172.31.100.1 255.255.255.0 no shut ! router ospf 100 vrf A net 172.31.200.1 0.0.0.255 area 0 A2 interface fa0/0 description To->ISP ip vrf forwarding A ip add 172.31.100.1 255.255.255.0 no shut ! router ospf 100 vrf A net 172.31.200.1 0.0.0.255 area 0

C. A1 interface fa0/0 description To->ISP ip vrf forwarding A ip add 172.31.100.1 255.255.255.0 no shut ! router ospf 100 net 172.31.100.1 0.0.0.255 area 0 A2 interface fa0/0 description To->ISP ip vrf forwarding A ip add 172.31.200.1 255.255.255.0 no shut ! router ospf 100 net 172.31.200.1 0.0.0.255 area 0

D. A1 interface fa0/0 description To->ISP ip add 172.31.200.1 255.255.255.0 no shut ! router ospf 100 net 172.31.200.1 0.0.0.255 area 0 A2 interface fa0/0 description To->ISP ip add 172.31.100.1 255.255.255.0 no shut ! router ospf 100 net 172.31.100.1 0.0.0.255 area 0

Correct Answer: *A*

Community vote distribution

A (83%) Other

Rui123 Highly Voted 1 year, 7 months ago

Selected Answer: A

Correct answer is A. Please note that A1, A2, B1 and B2 are Customer routers, therefore they have no idea of VRF.

upvoted 12 times

chris110 Most Recent 3 months, 3 weeks ago

Its A:
A1 interface fa0/0 description To->ISP
ip add 172.31.100.1 255.255.255.0
no shut
!
router ospf 100 net 172.31.100.1 0.0.0.255 area 0

A2 interface fa0/0 description To->ISP
ip add 172.31.200.1 255.255.255.0
no shut
!
router ospf 100 net 172.31.200.1 0.0.0.255 area 0

upvoted 2 times

HungarianDish 6 months, 4 weeks ago

Selected Answer: A

Answer "A", because we need to perform the configuration on the customer edges, and not on the ISP.
upvoted 3 times

AlexInShort12 3 days, 9 hours ago
The sentence is not really clear. After reading the question to fast, I though it was the other way around...
upvoted 1 times

Dacusai 8 months ago
Answer B and C are very similar but answer B has wrong IP address on router A2, answer C has the correct configuration for VRF lite to work.
upvoted 1 times

Koume 11 months, 1 week ago
Selected Answer: A
As the stament says VRF is fully configured on ISP, customerr routers just have to do is the necesary routing to communicate each other.
upvoted 1 times

Hurk2 11 months, 2 weeks ago
Selected Answer: A
A is correct, the ISP configures the VRF-lite not the customer
upvoted 1 times

PimplePooper 12 months ago
Selected Answer: A
A is the correct answer. Firstly, VRF is not configured on the client side that is already completed on the ISP router. Secondly, in the available answers none mentioned the creation of the VRF.
upvoted 1 times

stratosph3re 1 year ago
Selected Answer: A
Hello. The question clearly mentions that "The ISP router is FULLY CONFIGURED ... " , not the customer routers... Therefore, the correct answer is A because it doesn't contain any VRF info (Customers are unaware of the VRF concept) , and it has the correct IPs. Other than that, all the "vrf-related" answers are incorrect, because they don't mention anywhere the VRF-related creation commands .
upvoted 1 times

NoUserName1234 1 year, 1 month ago
Selected Answer: B
The question stated the they need too make use of the VRF lite Feature what implies that VRF Forwarding needs too be used. There is also a type error in the answers

See : https://itexamanswers.net/question/refer-to-the-exhibit-the-isp-router-is-fully-configured-for-customer-a-and-customer-b-using-the-vrf-lite-feature-what-is-the-minimum-configuration-required-for-customer-a-to-communicate-between-rout
upvoted 1 times

Koume 11 months, 1 week ago
What the question says on ISP vrf is fully configured. that means customer routers are not aware and have not to configue any vrf for this scenario to work.
upvoted 1 times

NoUserName1234 1 year, 1 month ago
The question stated the they need too make use of the VRF lite Feature what implies that VRF Forwarding needs too be used. There is also a type error in the answers

See : https://itexamanswers.net/question/refer-to-the-exhibit-the-isp-router-is-fully-configured-for-customer-a-and-customer-b-using-the-vrf-lite-feature-what-is-the-minimum-configuration-required-for-customer-a-to-communicate-between-rout
upvoted 1 times

Nhan 1 year, 5 months ago
It's clearly that A is correct answer, after assign vrf you setup ospf for routing in the vrf.
upvoted 1 times

piojo 1 year, 6 months ago
Selected Answer: A
A1 and A2 are CE, no VRFs there.
upvoted 3 times

tefacert 1 year, 7 months ago
i agree with A, since A1 and A2 are CE they dont need vrf config, its only necesary in Pe, in this case in ISP router
upvoted 4 times

Iarn 1 year, 7 months ago
Selected Answer: C
This is the only config which would work

```
A1
interface fa0/0
```

```
description To->ISP
ip vrf forwarding A
ip add 172.31.100.1 255.255.255.0
no shut
!
router ospf 100 net 172.31.100.1 0.0.0.255 area 0
A2
interface fa0/0
description To->ISP
ip vrf forwarding A
ip add 172.31.200.1 255.255.255.0
no shut
!
router ospf 100 net 172.31.200.1 0.0.0.255 area 0
```

upvoted 1 times

jthompaf 1 year, 7 months ago

I could be looking at this incorrectly, but if the ISP is fully configured with the proper VRFs on the interfaces connected to A1 and A2, no ip vrf forwarding command is necessary on A1 or A2. All they need is the right IP and Router OSPF with matching areas. With that being said, choice A works with the least amount of effort.

upvoted 3 times

Koume 11 months, 1 week ago

VRF on customer rotuers have not any sense if vrf is fully configured on ISP. this is the real use case scenario of VRF. isolate customers routing tables. Lab scenario A and you will notice.

upvoted 2 times

xziomal9 1 year, 7 months ago

```
C.
A1
interface fa0/0
description To->ISP
ip vrf forwarding A
ip add 172.31.100.1 255.255.255.0
no shut
!
router ospf 100 net 172.31.100.1 0.0.0.255 area 0
A2
interface fa0/0
description To->ISP
ip vrf forwarding A
ip add 172.31.200.1 255.255.255.0
no shut
!
router ospf 100 net 172.31.200.1 0.0.0.255 area 0

D.
A1
interface fa0/0
description To->ISP
ip add 172.31.200.1 255.255.255.0
no shut
!
router ospf 100
net 172.31.200.1 0.0.0.255 area 0
A2
interface fa0/0
description To->ISP ip add 172.31.100.1 255.255.255.0
no shut
!
router ospf 100 net 172.31.100.1 0.0.0.255 area 0
```

upvoted 3 times

xziomal9 1 year, 7 months ago

Selected Answer: A

```
A.
A1 interface fa0/0
description To->ISP
ip add 172.31.100.1 255.255.255.0
no shut
!
router ospf 100
net 172.31.100.1 0.0.0.255 area 0
A2
interface fa0/0
description To->ISP
ip add 172.31.200.1 255.255.255.0
no shut
!
router ospf 100 net 172.31.200.1 0.0.0.255 area 0
```

B.

```
A1 interface fa0/0
description To->ISP
ip vrf forwarding A
ip add 172.31.100.1 255.255.255.0
no shut
!
router ospf 100 vrf A
net 172.31.200.1 0.0.0.255 area 0
A2
interface fa0/0
description To->ISP
ip vrf forwarding A
ip add 172.31.100.1 255.255.255.0
no shut
!
router ospf 100 vrf A
net 172.31.200.1 0.0.0.255 area 0
```

upvoted 2 times

Elikem2 1 year, 8 months ago

Selected Answer: B

B. is the correct answer.
You have to configure "ip vrf forwarding A" under the interface before assigning IP address to the interface. If you do the reverse, the IP address will be removed.
You also have to specify the vrf in router ospf, i.e: "router ospf 100 vrf A"

upvoted 3 times

Question #73 *Topic 1*

An engineer is implementing a coordinated change with a server team. As part of the change, the engineer must configure interface GigabitEthernet2 in an existing VRF "RED" then move the interface to an existing VRF "BLUE" when the server team is ready. The engineer configured interface GigabitEthernet2 in VRF
"RED":
interface GigabitEthernet2
description Migration ID: B410A82D0935G35
vrf forwarding RED
ip address 10.0.0.0 255.255.255.254
negotiation auto
Which configuration completes the change?

A. interface GigabitEthernet2 no vrf forwarding RED vrf forwarding BLUE ip address 10.0.0.0 255.255.255.254

B. interface GigabitEthernet2 no ip address vrf forwarding BLUE

C. interface GigabitEthernet2 no vrf forwarding RED vrf forwarding BLUE

D. interface GigabitEthernet2 no ip address ip address 10.0.0.0 255.255.255.254 vrf forwarding BLUE

Correct Answer: *A*

Community vote distribution

A (100%)

MasterMatt 8 months, 2 weeks ago

Selected Answer: A

Answer A is correct but the IP address should be a suitable host IP and not a network id.

upvoted 1 times

Hurk2 11 months, 2 weeks ago

Selected Answer: A

A is correct, when removing a VRF or changing it, it will remove the IP address from the interface so the IP MASK will always need to be put after (conf-inf)vrf forwarding bla

upvoted 1 times

Nhan 1 year, 3 months ago

No vrf forwarding is needed, after you take the vrf off the int the IP address will be lost therefore you must reassign the up address

upvoted 1 times

piojo 1 year, 6 months ago

Selected Answer: A

Although "no vrf forwarding RED" is not needed

upvoted 2 times

Question #74 *Topic 1*

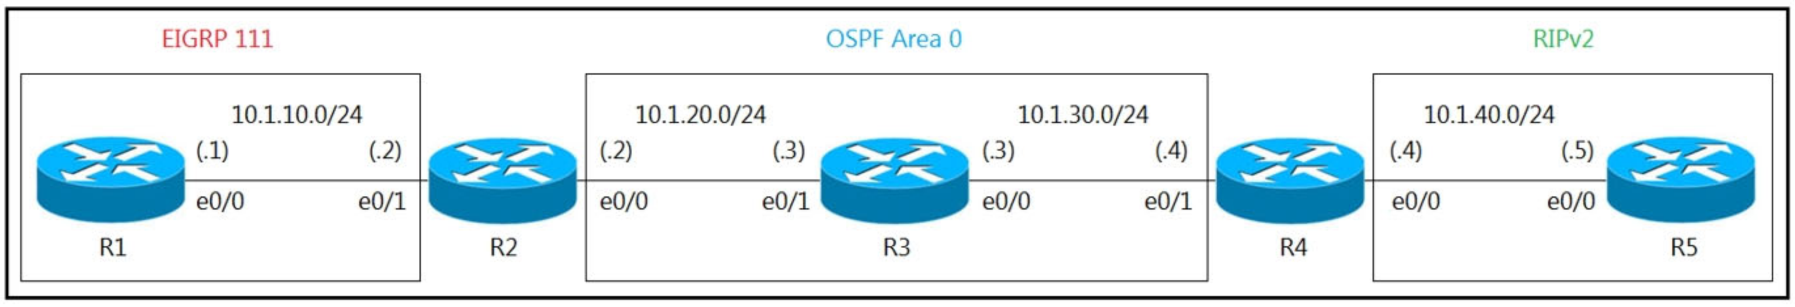

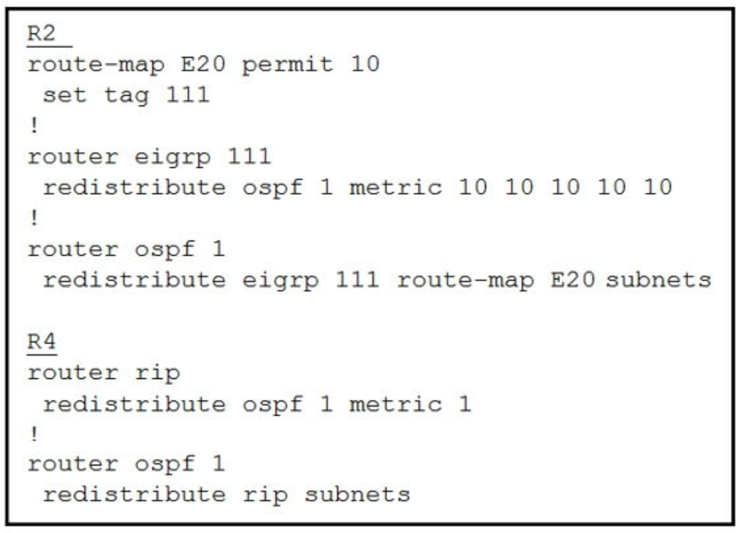

```
R2
route-map E20 permit 10
 set tag 111
!
router eigrp 111
 redistribute ospf 1 metric 10 10 10 10 10
!
router ospf 1
 redistribute eigrp 111 route-map E20 subnets

R4
router rip
 redistribute ospf 1 metric 1
!
router ospf 1
 redistribute rip subnets
```

Refer to the exhibit. R5 should not receive any routes originated in the EIGRP domain. Which set of configuration changes removes the EIGRP routes from the R5 routing table to fix the issue?

A. R4 route-map O2R deny 10 match tag 111 route-map O2R permit 20 ! router rip redistribute ospf 1 route-map O2R metric 1

B. R2 route-map E20 deny 20 R4 route-map O2R deny 10 match tag 111 ! router rip redistribute ospf 1 route-map O2R metric 1

C. R4 route-map O2R permit 10 match tag 111 route-map O2R deny 20 ! router rip redistribute ospf 1 route-map O2R metric 1

D. R4 route-map O2R deny 10 match tag 111 ! router rip redistribute ospf 1 route-map O2R metric 1

Correct Answer: *A*

Community vote distribution

A (100%)

AinsB 7 months ago

Selected Answer: A

Permit always important at the end of a route map to allow other routes to flow

upvoted 2 times

heeeeyajoke 1 year ago

Chosen answer is correct
R4
route-map O2R deny 10
match tag 111

route-map O2R permit 20 !

router rip
redistribute ospf 1 route-map O2R metric 1

upvoted 2 times

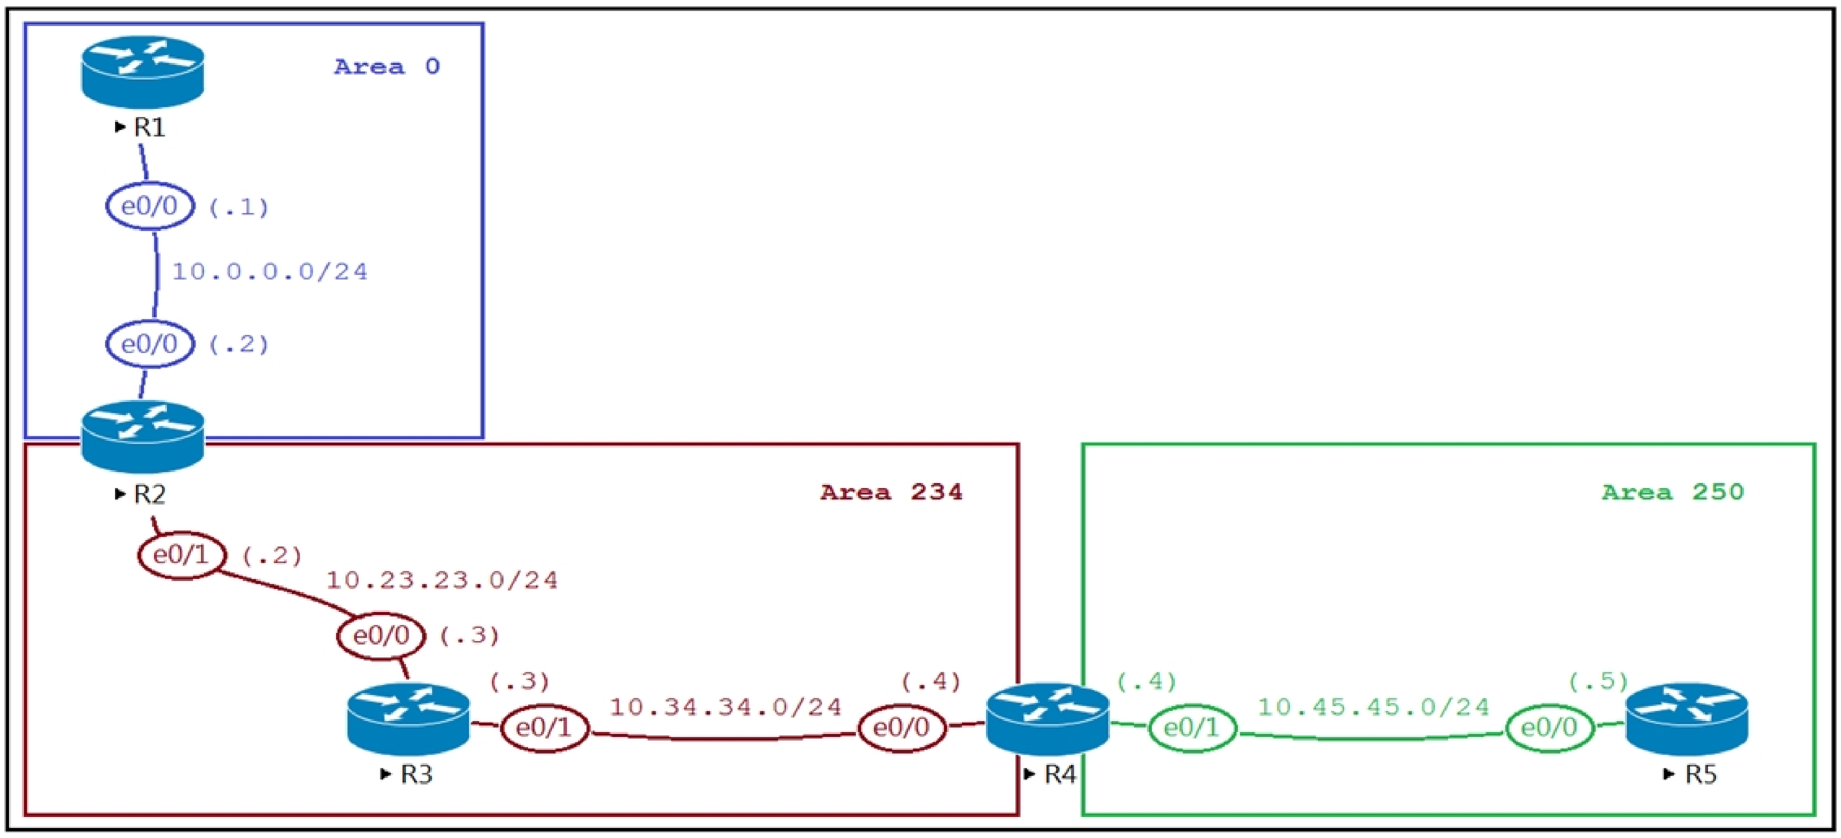

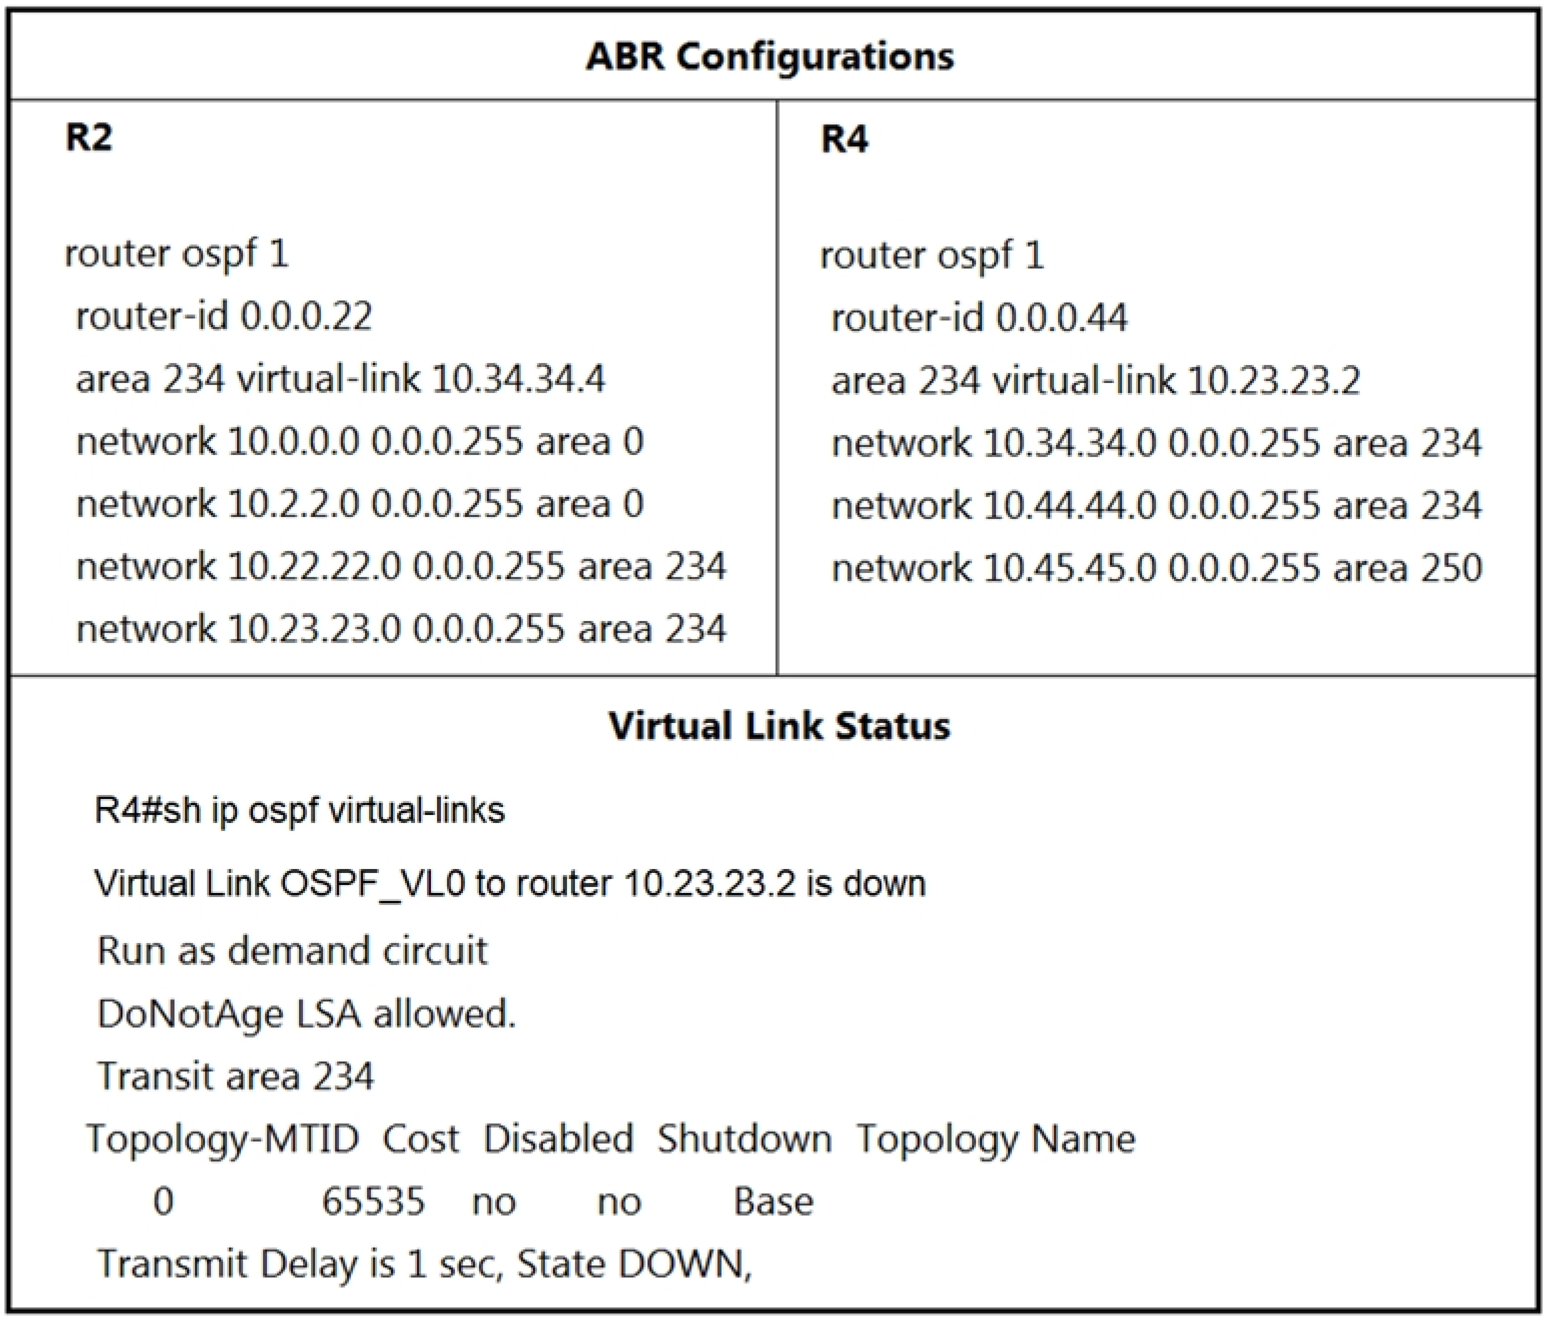

ABR Configurations

R2	R4
router ospf 1 router-id 0.0.0.22 area 234 virtual-link 10.34.34.4 network 10.0.0.0 0.0.0.255 area 0 network 10.2.2.0 0.0.0.255 area 0 network 10.22.22.0 0.0.0.255 area 234 network 10.23.23.0 0.0.0.255 area 234	router ospf 1 router-id 0.0.0.44 area 234 virtual-link 10.23.23.2 network 10.34.34.0 0.0.0.255 area 234 network 10.44.44.0 0.0.0.255 area 234 network 10.45.45.0 0.0.0.255 area 250

Virtual Link Status

```
R4#sh ip ospf virtual-links
Virtual Link OSPF_VL0 to router 10.23.23.2 is down
Run as demand circuit
DoNotAge LSA allowed.
Transit area 234
Topology-MTID  Cost  Disabled  Shutdown  Topology Name
      0          65535   no        no        Base
Transmit Delay is 1 sec, State DOWN,
```

Refer to the exhibit. The network administrator configured the network to connect two disjointed networks and all the connectivity is up except the virtual link, which causes area 250 to be unreachable.

Which two configurations resolve this issue? (Choose two.)

A. R2 router ospf 1 no area 234 virtual-link 10.34.34.4 area 234 virtual-link 0.0.0.44

B. R2 router ospf 1 no area 234 virtual-link 10.34.34.4 area 0 virtual-link 0.0.0.44

C. R4 router ospf 1 no area 234 virtual-link 10.23.23.2 area 0 virtual-link 0.0.0.22

D. R2 router ospf 1 router-id 10.23.23.2

E. R4 router ospf 1 no area 234 virtual-link 10.23.23.2 area 234 virtual-link 0.0.0.22

Correct Answer: *AE*

Reference:

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13703-8.html

Community vote distribution

AE (82%) BC (18%)

ZamanR 5 days, 1 hour ago

AE
Reference: https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13703-8.html

An important thing to remember when configuring virtual-link is we need to configure the OSPF router

ID and NOT the IP address of the ABR. Therefore in this question we have to use the command "area

234 virtual-link 0.0.0.44" on R2 and "area 234 virtual-link 0.0.0.22" on R4.

upvoted 1 times

chris110 3 months, 1 week ago

Question is wrong, the options are not like that in the exam.

upvoted 1 times

sajjad_gayyem 5 months, 3 weeks ago

Selected Answer: AE

A & E, check this
https://networklessons.com/ospf/how-to-configure-ospf-virtual-link

upvoted 1 times

HungarianDish 6 months, 4 weeks ago

Selected Answer: AE

For those, who see a virtual-link configuration with the router-id for the first time, like me:
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13703-8.html

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

Sorry, it was some misunderstanding by me, virtual-link is always configured with the OSPF router ID of the other ABR, and not the IP.
https://networklessons.com/ospf/how-to-configure-ospf-virtual-link

upvoted 2 times

Dominik_Networker 9 months, 3 weeks ago

Why not AD? R2 would be looking for the 0.0.0.44, what would be the ID of R4 and R4 would be looking for 10.23.23.2, what would be the ID of R2. Am I missing something?

upvoted 1 times

Koume 11 months, 1 week ago

Selected Answer: AE

The only error in the provided config is that were using the interface address instead of the router id

upvoted 1 times

Alexloh 1 year ago

Selected Answer: AE

AE is correct because each router points to the router ID of the other router.

upvoted 1 times

Mystic13 1 year, 7 months ago

Selected Answer: AE

AE are correct. BC are incorrect. The virtual link will use the transit area (area 234) to reach back to area 0 via router 0.0.0.22

upvoted 2 times

Iarn 1 year, 7 months ago

Selected Answer: AE

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/8313-27.html

upvoted 2 times

xziomal9 1 year, 7 months ago

Selected Answer: AE

The correct answer is: AE

upvoted 1 times

Kimaf 1 year, 8 months ago

Selected Answer: BC

A & E are wrong answers
B & C are the right answers as they use area 0 for virtual links to work

upvoted 2 times

⊟ **JOKERR** 1 year, 6 months ago

In the configuration of Virtual Links, we use the area of the trasnsit, not the backbone area. So A and E are correct.

upvoted 4 times

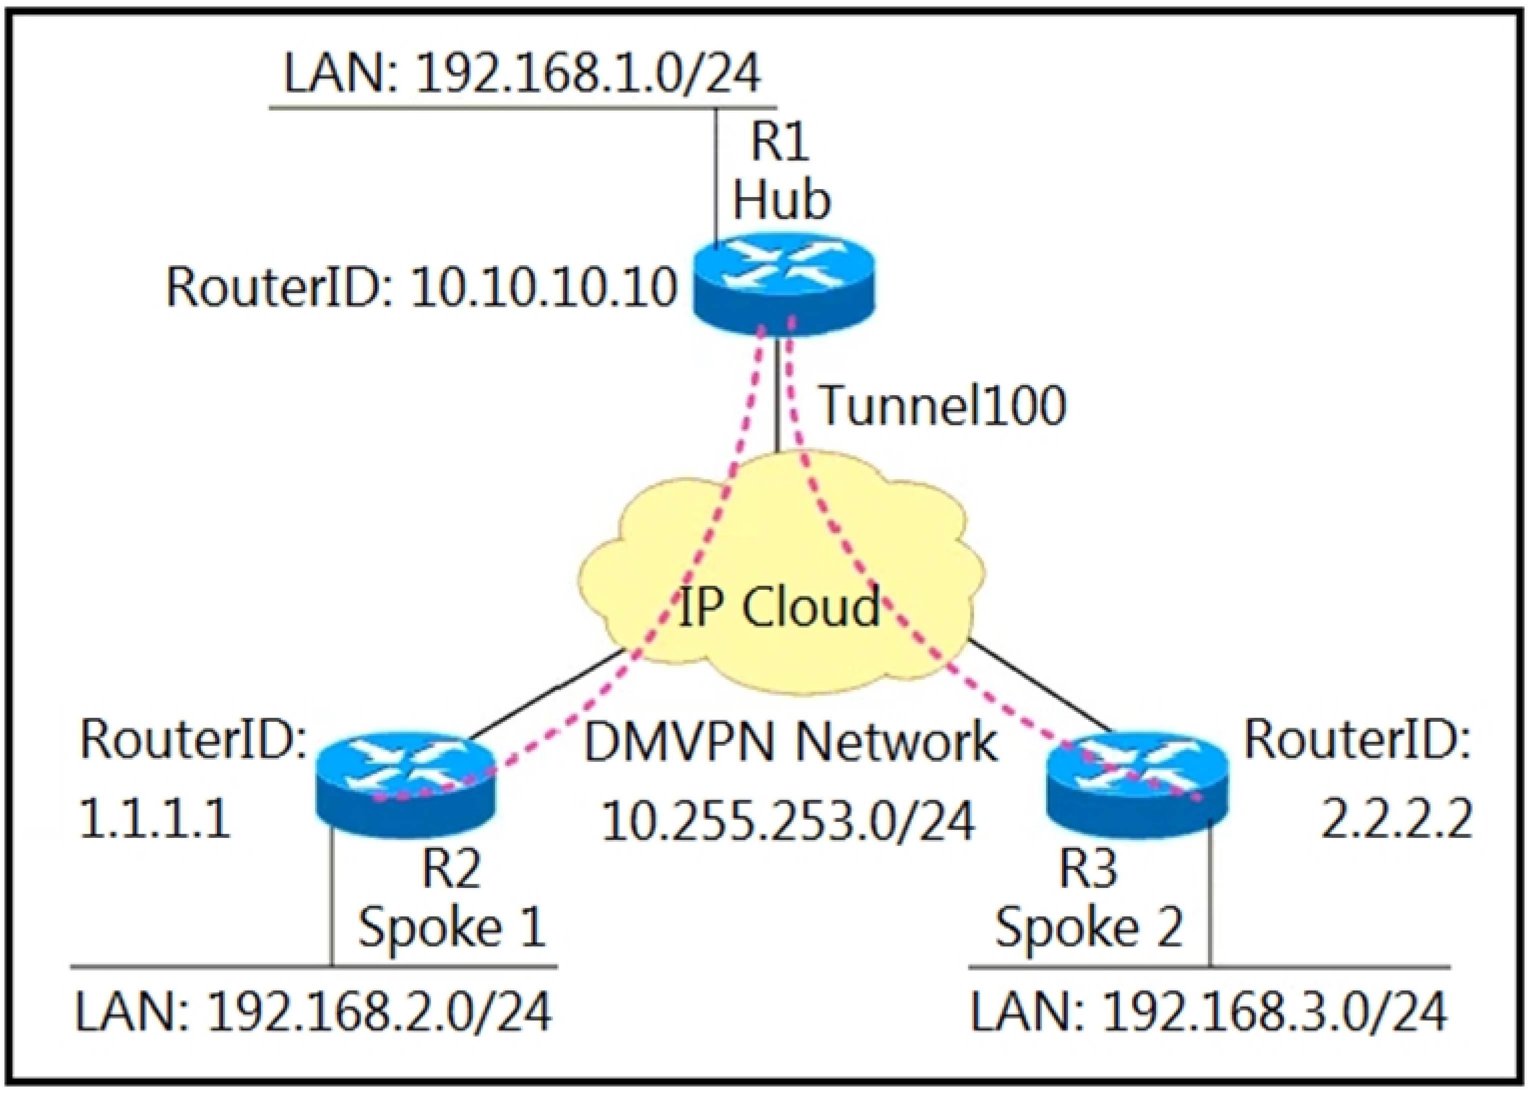

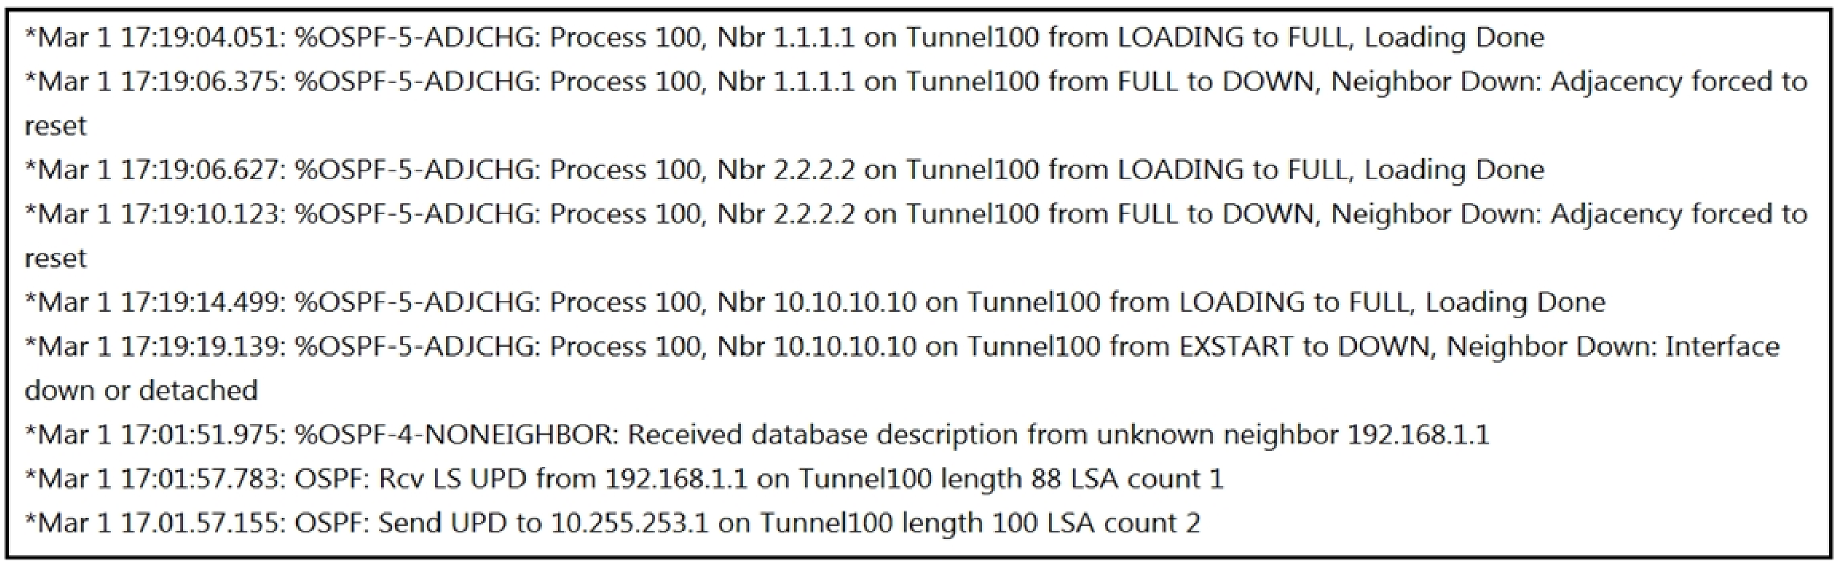

```
*Mar 1 17:19:04.051: %OSPF-5-ADJCHG: Process 100, Nbr 1.1.1.1 on Tunnel100 from LOADING to FULL, Loading Done
*Mar 1 17:19:06.375: %OSPF-5-ADJCHG: Process 100, Nbr 1.1.1.1 on Tunnel100 from FULL to DOWN, Neighbor Down: Adjacency forced to
reset
*Mar 1 17:19:06.627: %OSPF-5-ADJCHG: Process 100, Nbr 2.2.2.2 on Tunnel100 from LOADING to FULL, Loading Done
*Mar 1 17:19:10.123: %OSPF-5-ADJCHG: Process 100, Nbr 2.2.2.2 on Tunnel100 from FULL to DOWN, Neighbor Down: Adjacency forced to
reset
*Mar 1 17:19:14.499: %OSPF-5-ADJCHG: Process 100, Nbr 10.10.10.10 on Tunnel100 from LOADING to FULL, Loading Done
*Mar 1 17:19:19.139: %OSPF-5-ADJCHG: Process 100, Nbr 10.10.10.10 on Tunnel100 from EXSTART to DOWN, Neighbor Down: Interface
down or detached
*Mar 1 17:01:51.975: %OSPF-4-NONEIGHBOR: Received database description from unknown neighbor 192.168.1.1
*Mar 1 17:01:57.783: OSPF: Rcv LS UPD from 192.168.1.1 on Tunnel100 length 88 LSA count 1
*Mar 1 17.01.57.155: OSPF: Send UPD to 10.255.253.1 on Tunnel100 length 100 LSA count 2
```

Refer to the exhibit. A network administrator sets up an OSPF routing protocol for a DMVPN network on the hub router.
Which configuration command is required to establish a DMVPN tunnel with multiple spokes?

A. ip ospf network point-to-point on the hub router

B. ip ospf network point-to-multipoint on one spoke router

C. ip ospf network point-to-multipoint on both spoke routers

D. ip ospf network point-to-point on both spoke routers

Correct Answer: *C*

Community vote distribution

C (64%) D (36%)

louisvuitton12 1 month, 3 weeks ago

Selected Answer: D

The correct answer is D. Look at the purple lines, this is Hub to Spoke, this is NOT Spoke to Spoke topology. So Option C is definitely not the answer.

upvoted 1 times

no_name995 4 months, 3 weeks ago

Answer C: this linked helped a lot - https://community.cisco.com/t5/routing/help-with-ospf-and-dmvpn/td-p/2107338

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: C

we have two neigbors, beacuse is necesary point-to-multipoint

upvoted 1 times

HungarianDish 8 months ago

Selected Answer: C

This is how I see it: The issue "Tunnel100 from EXCHANGE to DOWN, Neighbor Down: Adjacency forced to reset" is caused by OSPF default network type p2p.
There is a different solution per DMVPN Phase. OSPF network type broadcast is suitable for phase 1 and 2, whereas point-to-multipoint is suitable for Phase3.
We can exclude answers with p2p, like A and D. The solution needs to be applied on all spokes, so answer C is fitting best.

https://community.cisco.com/t5/routing/help-with-ospf-and-dmvpn/td-p/2107338/page/2
https://networklessons.com/cisco/ccie-routing-switching/dmvpn-phase-1-ospf-routing

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

https://networklessons.com/cisco/ccie-routing-switching/dmvpn-phase-3-ospf-routing#Point-to-multipoint

upvoted 1 times

Koume 11 months, 1 week ago

Selected Answer: C

I lab the given anwer is correct spoke be point to multpoint to that scenario work

upvoted 3 times

Noproblem22 1 year, 1 month ago

I believe C is correct

upvoted 2 times

wts 1 year, 3 months ago

Selected Answer: C

Ahh, again nothing is clear.

So. Obviously, the problems are connected EXACTLY with the wrong type of network. I will assume that a point-to-point is configured, and the router needs to establish a neighborhood with two on the same interface, so it blinks(point-to-2point).
Second. Why only on one spoke(is this a typical and highly scalable role for a router)? Of course for everyone.
According to my logic, it turns out C.

P.S.: Does anyone understand how to determine the phase here?

upvoted 1 times

Koume 11 months, 1 week ago

When i did the lab of that scenario, by default tunnel interfaces are point to point, leaving by default osp start flapping adjacencies, the hub by desing must be point-to-multipoint, as gre interface also is multipoint. the spoke as are p2p by default on GRE interfaces then start failing as the hello and dead timer of point-to-multipoint and point-to-point are differnt. So the mos correst answer is C as you state.

upvoted 1 times

Koume 11 months, 1 week ago

I go here for D, as the point to multpoint is only significant on HUB router. but this in phase 2 on phase 3 also the spokes shoulbe point to multipoint . As the question is pretty ambiguos based on the answers.

upvoted 2 times

Edwinmolinab 1 year, 4 months ago

Selected Answer: C

given answer is correct according to https://www.grandmetric.com/knowledge-base/design_and_configure/dmvpn-phase-3-single-hub-ospf-spoke-example/

upvoted 1 times

WAKIDI 1 year, 6 months ago

Selected Answer: D

the right answer is D. ip ospf network point-to-point on both spoke routers
reference : https://learningspace.cisco.com/ Book title : ENARSI - Implementing Cisco Enterprise Advanced Routing and Services - Student Learning Guide, version=1.0.20, page 215
: "In strict hub-and-spoke DMVPNs, you should include the tunnel interface in the OSPF routing process, and configure the tunnel interface as a point-to-multipoint OSPF network type on the hub router, and as a point-to-point network type on the branch routers. In this case, there is no need to elect a DR on the DMVPN subnet."

upvoted 3 times

TECH3K3 1 year, 5 months ago

Also OSPF default is point-to-point and we are using multipoint interfaces for each spoke to see each other

upvoted 1 times

dongzh007 1 year, 5 months ago

D is wrong.
ospf use dmvpn phase3, all hub and spokes should be point-to-multipoint.

upvoted 1 times

⊟ **abd123** 10 months, 2 weeks ago

who did you know that is using phase 3

upvoted 1 times

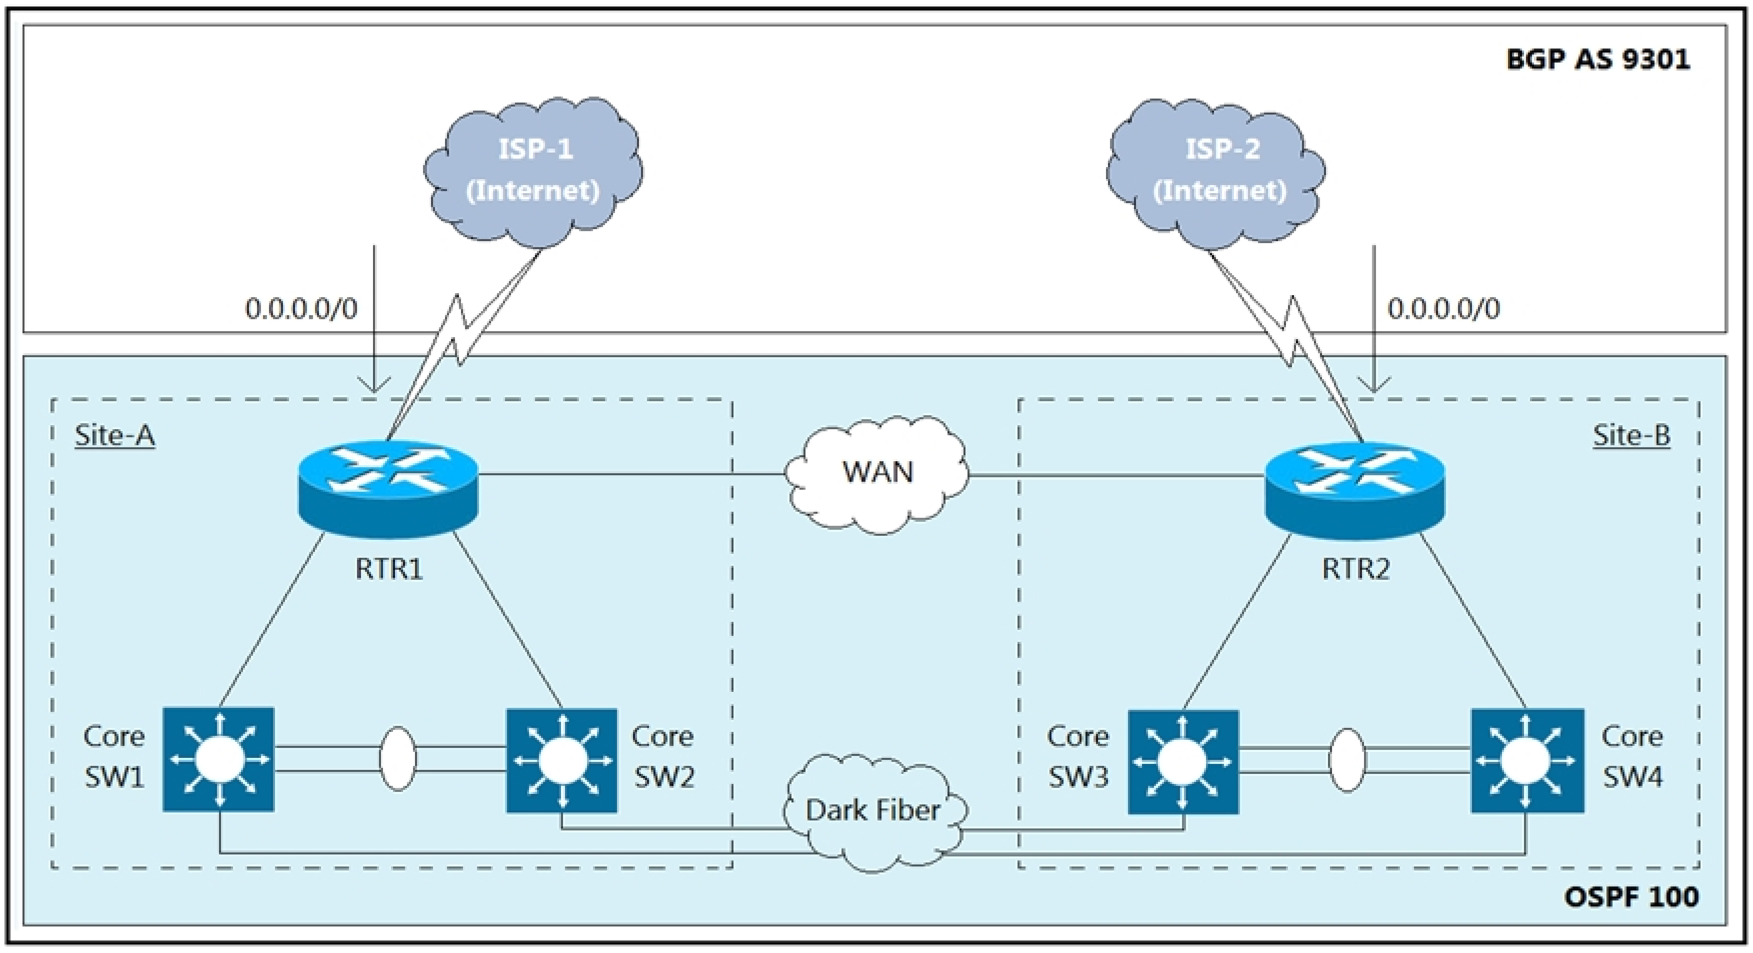

Refer to the exhibit. The Internet traffic should always prefer Site-A ISP-1 if the link and BGP connection are up; otherwise, all Internet traffic should go to ISP-2.

Redistribution is configured between BGP and OSPF routing protocols, and it is not working as expected.

What action resolves the issue?

A. Set OSPF Cost 200 at Site-A RTR1, and set OSPF Cost 100 at Site-B RTR2.

B. Set metric-type 2 at Site-A RTR1, and set metric-type 1 at Site-B RTR2.

C. Set metric-type 1 at Site-A RTR1, and set metric-type 2 at Site-B RTR2.

D. Set OSPF Cost 100 at Site-A RTR1, and set OSPF Cost 200 at Site-B RTR2.

Correct Answer: *C*

Community vote distribution

C (92%) 8%

piojo Highly Voted 1 year, 6 months ago

Selected Answer: C

OSPF prefers E1 over E2

upvoted 5 times

louisvuitton12 Most Recent 1 month, 3 weeks ago

Selected Answer: C

The order of preference for OSPF as per RFC 2328 is :

intra-area routes, O
interarea routes, O IA
external routes type 1, O E1
external routes type 2, O E2
This rule of preference cannot be changed.

upvoted 1 times

[Removed] 4 months, 2 weeks ago

Selected Answer: C

As already explained. Type 1 over type 2, and to add to the explanation. Remember that an external route redistributed into OSPF is by default type 2, this is why the initial design was not working as intended. The default route from ISP1 needed to be defined as type 1

upvoted 1 times

HungarianDish 6 months, 4 weeks ago

Selected Answer: C

https://networklessons.com/ospf/ospf-path-selection-explained
OSPF will first look at the “type of path” to make a decision and, secondly look at the metric (cost).

type of path - path selection order: O > O IA > N1 > E1 > N2 > E2
upvoted 1 times

AinsB 7 months, 1 week ago

Selected Answer: D

Before it is redistributed you need the IGP to prefer a path, in this case changing the ospf cost will allow one path to be preferred over the other
upvoted 1 times

baldebri 11 months, 2 weeks ago

setting the costs 100 and 200 and not calculating it will make sit-A preferred immediately on the routers regardless of adding bandwidth to the destination D is correct
upvoted 1 times

Orchidium 1 year, 5 months ago

Selected Answer: C

User piojo is correct. Answer is C. E1 over E2 routes all day, regardless of cost.
upvoted 1 times

xziomal9 1 year, 7 months ago

Selected Answer: C

The correct answer is: C
upvoted 3 times

JOKERR 1 year, 6 months ago

Can you explain why? I believe setting OSPF cost to 100 at SITE1 would select that site to be the exit path according to BGP Best PAs...

Again, I am not sure. Just looking for an explanation.
upvoted 1 times

JingleJangus 1 year, 6 months ago

C is likely a better answer because modifying cost could require manually modifying several interfaces to get the intended effect... whereas just simply modifying external type during redistribution immediately allows all routers to prefer external type1 routes over external type2.
upvoted 4 times

Question #78

Topic 1

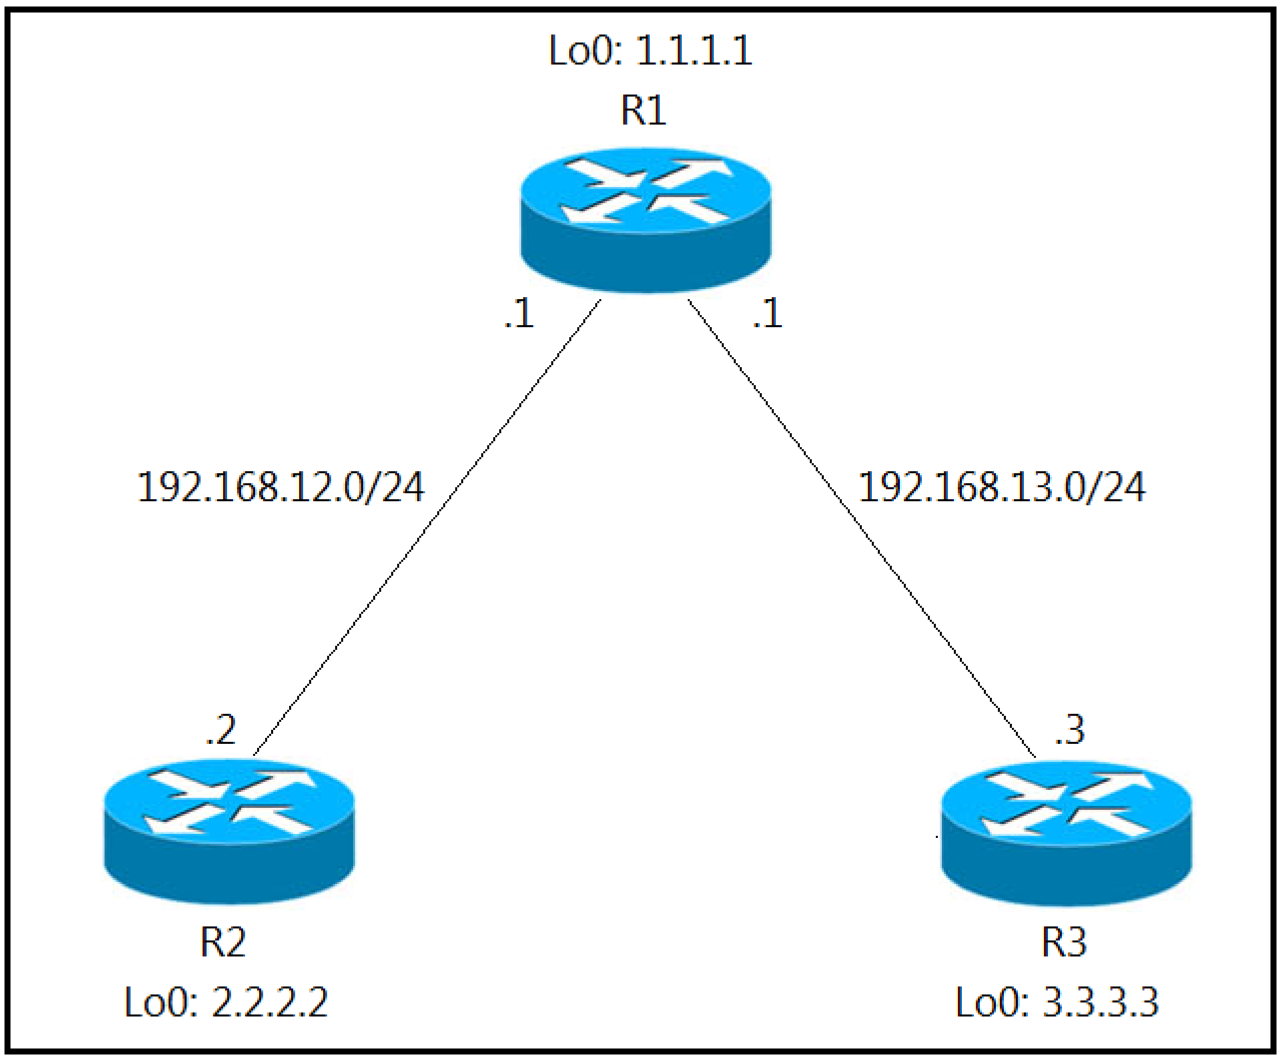

Refer to the exhibit. An engineer has configured R1 as EIGRP stub router. After the configuration, router R3 failed to reach to R2 loopback address.

Which action advertises R2 loopback back into the R3 routing table?

A. Add a static route for R2 loopback address in R1 and redistribute it to advertise to R3.

B. Use a leak map on R1 that matches the required prefix and apply it with the distribute list command toward R3.

C. Use a leak map on R3 that matches the required prefix and apply it with the EIGRP stub feature.

D. Add a static null route for R2 loopback address in R1 and redistribute it to advertise to R3.

Correct Answer: *B*

Community vote distribution

B (72%) A (22%) 6%

⊟ **louisvuitton12** 1 month, 3 weeks ago

Selected Answer: B

Option B is correct:
https://networklessons.com/cisco/ccie-routing-switching-written/eigrp-stub-leak-map

upvoted 2 times

⊟ **Pietjeplukgeluk** 1 month ago

The example you supplied matches the question. So i agree it should be B. Also strange that we get this question on CCNP as it seems out of scope of the exam. By the way, love the example, so simple!

upvoted 1 times

⊟ **jansan55** 3 months, 2 weeks ago

Selected Answer: A

As previously mentioned (for example by HungarianDish) there is no such possibilty eigrp leak-map with distribute list. I also labbed the configuration, and I agree with HungarianDish, that answer is "A". Even the wording is met.

upvoted 1 times

⊟ **Brand** 3 months, 2 weeks ago

Selected Answer: B

I don't know why but this questions seems like it's asking about leak-map feature of EIGRP stub configuration... I can't explain but I have this feeling maybe it's because there is no other question talking about leak-map. Therefor it's B to me.

upvoted 1 times

⊟ **inteldarvid** 4 months, 4 weeks ago

Selected Answer: A

the answer correct is "A"

upvoted 1 times

⊟ **inteldarvid** 4 months, 4 weeks ago

because, need redisitribute static in R1, because R1 is router stub

upvoted 1 times

⊟ **[Removed]** 5 months, 1 week ago

Selected Answer: B

A and B are correct, but both of them have terrible wording.

A.
If you configure a static route to 2.2.2.2, then that will remove that route from the EIGRP topology due to AD. So even if R1 was not a stub router, it would not advertise the route.
On top of that, we not only have to redistribute static routes into EIGRP, we also need to configure the eigrp stub to advertise static, as follows:

```
eigrp <AS#>
redistribute static <---This alone does not redistribute the static route
eigrp stub static <---This completes the static redistribute for a stub router
```

B
There is not distribute-list into eigrp stub command. But what I think cisco meant was to create a prefix-list for R2 Loopback, reference it into a route-map, and tie that to a leak-map in eigrp, as follows:

```
ip prefix-list R2LOOP permit 2.2.2.2/32
!
route-map R2LOOP permit
match ip address prefix-list R2LOOP
!
router eigrp <AS#>
eigrp stub leak-map R2LOOP
```

So I think B is the best answer, as A is far too vague.

upvoted 1 times

⊟ **adudeguy** 6 months, 1 week ago

Answer C.
A & D both require "redistribute static" under EIGRP. B is wrong because the leak-map is applied to the stub command and not via a distribute list.

upvoted 1 times

⊟ **Malasxd** 7 months, 2 weeks ago

The answer "A" would work, but "B" makes so much more sense for me.

https://networklessons.com/cisco/ccie-routing-switching-written/eigrp-stub-leak-map

upvoted 1 times

⊟ **HungarianDish** 6 months, 4 weeks ago

Hi! There seems to be a problem with answer "B". As already pointed out in other comments, the "eigrp stub leak-map" is using a route-map, and not a distribute-list.

upvoted 2 times

⊟ **HungarianDish** 8 months ago

Using the "redistributed" parameter had the very same effect:
Config on R1:

```
!
router eigrp 1
network 1.1.1.1 0.0.0.0
network 192.168.12.0
network 192.168.13.0
redistribute static
eigrp stub connected summary redistributed

ip route 2.2.2.2 255.255.255.255 192.168.12.2
!
```

upvoted 1 times

⊟ **HungarianDish** 8 months ago

For me, definitely answer "A".

upvoted 1 times

⊟ **HungarianDish** 8 months ago

After applying the config, that is what you see on R3:

```
r3#sh ip route eigrp | b Gateway
Gateway of last resort is not set
```

```
1.0.0.0/32 is subnetted, 1 subnets
D 1.1.1.1 [90/130816] via 192.168.13.1, 00:13:36, GigabitEthernet0/0
2.0.0.0/32 is subnetted, 1 subnets
D EX 2.2.2.2 [170/3072] via 192.168.13.1, 00:12:49, GigabitEthernet0/0
D 192.168.12.0/24 [90/3072] via 192.168.13.1, 00:13:36, GigabitEthernet0/0
r3#
```
upvoted 1 times

HungarianDish 8 months ago

My config on R1:
```
router eigrp 1
network 1.1.1.1 0.0.0.0
network 192.168.12.0
network 192.168.13.0
redistribute static
eigrp stub connected static summary

ip route 2.2.2.2 255.255.255.255 192.168.12.2
```

You need the static route on R1 to get the prefix 2.0.0.0/32 back into the eigrp table. Plus, you need to push out the static route with these two commands on R1:
```
redistribute static
eigrp stub static
```
upvoted 1 times

HungarianDish 8 months ago

*typo: prefix 2.2.2.2/32

upvoted 1 times

HungarianDish 8 months ago

Selected Answer: A

For me "A" is correct. I labbed it in CML. "B" is incorrect, because there is no such distribute-list (as others pointed out). Pls, see some good explanations here:

https://community.cisco.com/t5/routing/eigrp-eigrp-stub-connected-static-summary/td-p/2575321
https://learningnetwork.cisco.com/s/question/0D53i00000Kso40CAB/clarification-needed-on-eigrp-stub-static-option
https://notes.networklessons.com/redistribute-static-route-into-eigrp-stub-router

upvoted 2 times

zhlzjz 11 months, 3 weeks ago

Use a leak map on R1 that matches the required prefix and apply it with the EIGRP stub feature.

upvoted 2 times

Typovy 8 months, 3 weeks ago

Yea we know that but there is no answer like that so we have to choose something from cisco fuc*ed up answers. So answer is B

upvoted 1 times

zhlzjz 11 months, 3 weeks ago

I lab it
answer B is correct.
distribute-list works foradvise 2.2.2.2. It look like distribute-list is over eigrp stub feature.
If C is configure in R1, Also correct.

upvoted 4 times

TheBaja 1 year, 1 month ago

Answer is C. eigrp stub leak-map MY-ROUTE-MAP make exeption for stub feature.
```
router eigrp 1
eigrp stub leak-map MY-ROUTE-MAP
route-map MY-ROUTE-MAP permit 10
match ip address MY-ACL
ip access-list standard MY-ACL
permit x.x.x.x wildcard (loopback network)
```
upvoted 2 times

Noproblem22 1 year, 1 month ago

I believe B is correct https://networklessons.com/cisco/ccie-routing-switching-written/eigrp-summary-leak-map

upvoted 1 times

toto89 1 year, 1 month ago

Leak-map can't be applied to distribute-lists it can only be applied to the stub so B is incorrect.
A could be correct because R1 stub could be configured to advertise redistributed routes. In that case the route with null0 would be advertised to both routers. The route wouldn't become a successor on R2 because it already has a connected route with an AD of 1 which is lower than 90 so the R2 wouldn't add this route into his routing table. R3 on the other hand would add it to his routing table.

upvoted 1 times

JoeyT 6 months ago

who said you would apply distribution list TO leap-map? it says WITH it.

upvoted 1 times

Huntkey 1 year, 3 months ago

B says applied with a distribute-list and that is incorrect, isn't it?

upvoted 3 times

ntdevera 1 year, 4 months ago

Selected Answer: B

Its B
Redistributing static routes into eigrp isnt enough. If you don't have the "eigrp stub static" configured, it will not advertise the redistributed route.

Ref: https://learningnetwork.cisco.com/s/article/eigrp-stub-routing

upvoted 1 times

Question #79 *Topic 1*

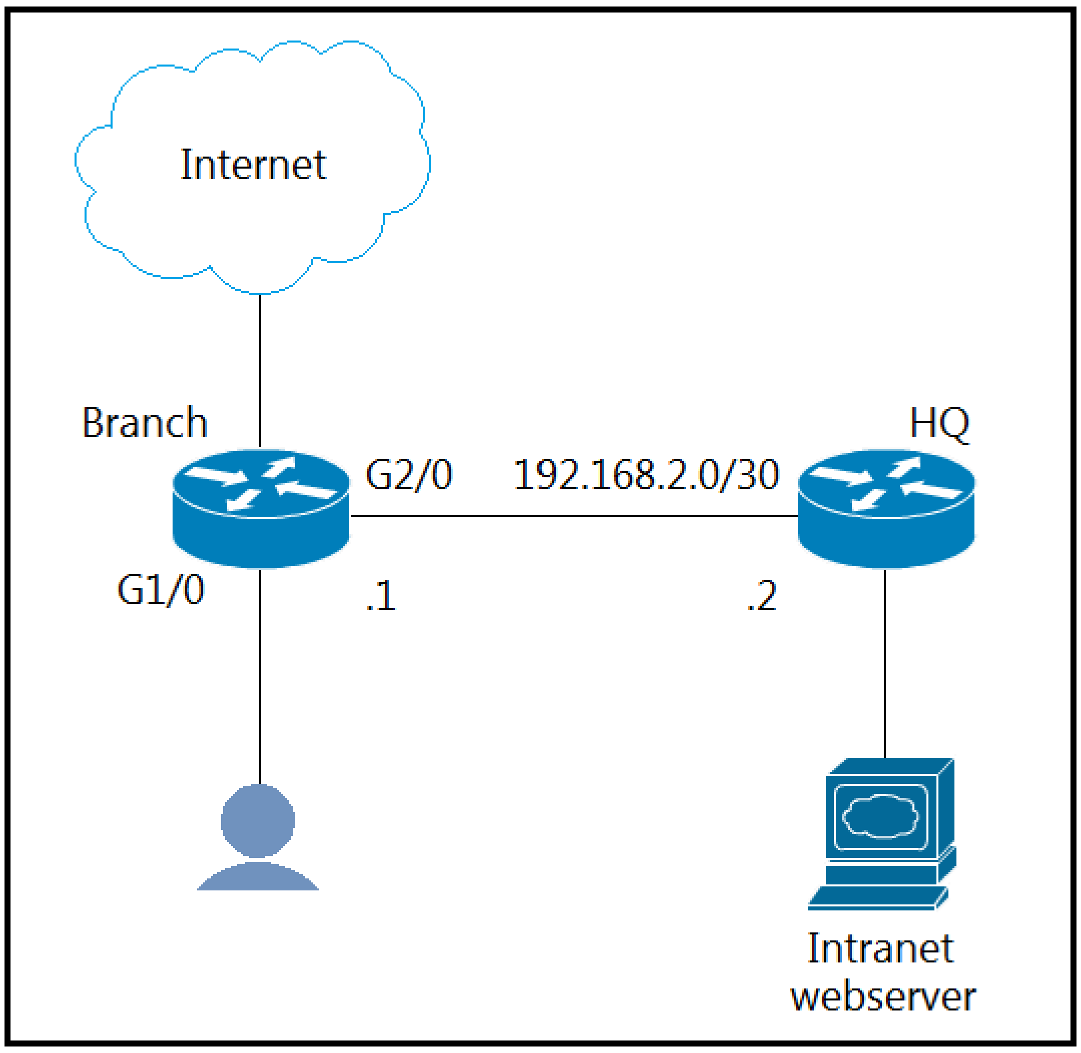

Refer to the exhibit. The branch router is configured with a default route toward the Internet and has no routes configured for the HQ site that is connected through interface G2/0. The HQ router is fully configured and does not require changes.
Which configuration on the branch router makes the intranet website (TCP port 80) available to the branch office users?

A. access-list 101 permit tcp any any eq 80 access-list 102 permit tcp any host intranet-webserver-ip ! route-map pbr permit 10 match ip address 101 set ip next-hop 192.168.2.2 route-map pbr permit 20 match ip address 102 set ip next-hop 192.168.2.2 ! interface G2/0 ip policy route-map pbr

B. access-list 100 permit tcp host intranet-webserver-ip eq 80 any ! route-map pbr permit 10 match ip address 100 set ip next-hop 192.168.2.2 ! interface G1/0 ip policy route-map pbr

C. access-list 100 permit tcp any host intranet-webserver-ip eq 80 ! route-map pbr permit 10 match ip address 100 set ip next-hop 192.168.2.2 ! interface G2/0 ip policy route-map pbr

D. access-list 101 permit tcp any any eq 80 access-list 102 permit tcp any host intranet-webserver-ip ! route-map pbr permit 10 match ip address 101 102 set ip next-hop 192.168.2.2 ! interface G1/0 ip policy route-map pbr

Correct Answer: *D*

Community vote distribution

D (75%) C (25%)

Cyril_the_Squirl 4 months, 1 week ago

By process of elimination (A) & (C) = PRB applied on wrong interface. (B) wrong ACL syntax, leaving D as the only right option.

upvoted 1 times

[Removed] 4 months, 1 week ago

Selected Answer: D

D, the instructions say that the intranet branch users require to have access to the intranet web server at HQ without modifying the routing table at Branch, the only way is to point all the Branch network users to the next hop 192.168.2.2 on TCP port 80. Therefore the PBR has to be applied at Branch router interface G1/0

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

the option correct is D, beacause PBR match with interface g1/0 (gateway user)

upvoted 1 times

Chiaretta 5 months, 3 weeks ago

Selected Answer: D

```
A:
access-list 101 permit tcp any any eq 80
access-list 102 permit tcp any host intranet-webserver-ip
route-map pbr permit 10 match ip address 101
set ip next-hop 192.168.2.2
route-map pbr permit 20 match ip address 102
set ip next-hop 192.168.2.2
interface G2/0 ip policy route-map pbr

B:
access-list 100 permit tcp host intranet-webserver-ip eq 80 any
route-map pbr permit 10 match ip address 100
set ip next-hop 192.168.2.2
interface G1/0 ip policy route-map pbr

C:
access-list 100 permit tcp any host intranet-webserver-ip eq 80
route-map pbr permit 10 match ip address 100
set ip next-hop 192.168.2.2
interface G2/0 ip policy route-map pbr

D:
access-list 101 permit tcp any any eq 80
access-list 102 permit tcp any host intranet-webserver-ip
route-map pbr permit 10 match ip address 101 102
set ip next-hop 192.168.2.2
interface G1/0 ip policy route-map pbr
```

PBR must be placed on traffic ingress interface.
upvoted 2 times

Dacusai 8 months ago

I don't see a correct answer here, you can not send all http traffic to the intranet server in this case, in this case C is more likely because it only will apply to traffic destinated to the server but is missing the permit 20 on the route map.
upvoted 3 times

Pietjeplukgeluk 1 month ago

Using Policy Based Routing there is no requirement for "route-map route_map_name permit 20" as in this case when no policy base routing is used, normal routing is used. So do not mix applying a route-map as route filtering (that has an implicit deny) and applying a route map for PBR. Anyway, in my opinion C is also correct, only it is applied to the wrong interface.
upvoted 1 times

HungarianDish 8 months ago

Selected Answer: D

"C" is for egress traffic, "D" is for ingress, so for me "D" is right.
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2/25ew/configuration/guide/conf/pbroute.pdf
"You specify PBR on the incoming interface (the interface on which packets are received), not outgoing interface."
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/15-0SY/configuration/guide/15_0_sy_swcg/policy_based_routing_pbr.pdf
"PBR cannot be applied to egress traffic or to multicast traffic."
upvoted 2 times

6dd4aa0 8 months, 3 weeks ago

Selected Answer: C

Answer C does the job accordingly to the question asked.

Answer D is more generally conditions which will work too.
upvoted 1 times

Titini 10 months ago

Selected Answer: D

I believe it is D as it is applied in the correct interface G1/0.
upvoted 2 times

Koume 11 months, 1 week ago

Selected Answer: D

To me seems the more right even if pass all 80 traffic to web server.
upvoted 1 times

rogabor81 11 months, 3 weeks ago

Selected Answer: D

The best answer would be C if the pbr is applied to Gi0/1 and not Gi0/2.
In the given answers D is the closest one, but it sends EVERY HTTP(port80) traffic sourced from Branch to the Intranet webserver. Considering that you probaply never want to allow your network to communicate through open HTTP(80) on the internet, this makes more sense then any other option.
upvoted 1 times

Alexloh 1 year ago

Selected Answer: C

Answer C looks more logical compared to D.

upvoted 1 times

Koume 11 months, 1 week ago

No, because on C is applying to the outbound interface GI0/2, so PBR will never match as PBR works when analizing the inbond interface.

upvoted 1 times

DUBC89x 1 year ago

C.

```
access-list 100
permit tcp any host intranet-webserver-ip eq 80
!
route-map pbr permit 10
match ip address 100
set ip next-hop 192.168.2.2
!
interface G2/0
ip policy route-map pbr
```

upvoted 2 times

CisconAWSGURU 1 year, 1 month ago

Selected Answer: C

C, makes sense to me!

upvoted 1 times

NoUserName1234 1 year, 1 month ago

All Answer are techically wrong Answer D makes all traffic flow to HQ instead of only the Web Traffic as stated in the qeustion. A is also wrong due too outgoing interface B is Fault in the syntax of the ACL Answer C is also outgoing interface

upvoted 3 times

jarz 1 year, 2 months ago

Selected Answer: C

You only need the single ACL to match the Internet webserver IP .

upvoted 2 times

jarz 1 year, 1 month ago

I actually retract my answer, none are correct.
D is the closest to being correct.

upvoted 1 times

babs 1 year, 2 months ago

the same job can be done via option B,

upvoted 1 times

Huntkey 1 year, 3 months ago

The only problem is that this would send all traffic including to the internet on port 80 to the hub router

upvoted 2 times

Reikidude00 1 year, 5 months ago

Selected Answer: D

Can anyone explain the meaning to configure route the internet traffic to HQ router?
Why access list 101 should be specified in the match address? it works but it is sub-optimal routing

upvoted 1 times

Question #80 Topic 1

R1 and R2 are configured as eBGP neighbors. R1 is in AS100 and R2 is in AS200. R2 is advertising these networks to R1:
172.16.16.0/20
172.16.3.0/24
172.16.4.0/24
192.168.1.0/24
192.168.2.0/24
172.16.0.0/16
The network administrator on R1 must improve convergence by blocking all subnets of 172.16.0.0/16 major network with a mask lower than 23 from coming in.
Which set of configurations accomplishes the task on R1?

A. ip prefix-list PL-1 deny 172.16.0.0/16 ge 23 ip prefix-list PL-1 permit 0.0.0.0/0 le 32 ! router bgp 100 neighbor 192.168.100.2 remote-as 200 neighbor 192.168.100.2 prefix-list PL-1 in

B. ip prefix-list PL-1 deny 172.16.0.0/16 le 23 ip prefix-list PL-1 permit 0.0.0.0/0 le 32 ! router bgp 100 neighbor 192.168.100.2 remote-as 200 neighbor 192.168.100.2 prefix-list PL-1 in

C. ip prefix-list PL-1 deny 172.16.0.0/16 ip prefix-list PL-1 permit 0.0.0.0/0 ! router bgp 100 neighbor 192.168.100.2 remote-as 200 neighbor 192.168.100.2 prefix-list PL-1 in

D. access-list 1 deny 172.16.0.0 0.0.254.255 access-list 1 permit any ! router bgp 100 neighbor 192.168.100.2 remote-as 200 neighbor 192.168.100.2 distribute-list 1 in

Correct Answer: *B*

Community vote distribution

A (74%) B (26%)

Cyril_the_Squirl Highly Voted 4 months, 1 week ago

It looks like nobody has read the question :-)
The answer is B

upvoted 7 times

Pietjeplukgeluk 1 month ago

If actually agree here, blocking the "less specific" routes also reduces advertised routes. And the " mask lower than 23" is clearly stating 23 and lower. As the question is stupid, i agree, and anyone picking A has a point, it makes more sense, but anyway, it is not the question.

upvoted 1 times

kaupz 6 days, 18 hours ago

a mask lower than 23 - this means mask 22, 21, 20 ... 16 - I would go for B. But ofcourse IRL you would do the other way around.

upvoted 1 times

Slinky Highly Voted 8 months, 3 weeks ago

Selected Answer: A

Answer is A. Question is worded awfully, but seeing as they want you to "improve convergence" that would imply you are trying to reduce the amount of prefixes that you are accepting from AS200. Only ge 23 is going to get you there.

upvoted 6 times

net_eng10021 3 months, 3 weeks ago

Slinky nails it....'improve convergence', hence the goal is to reduce table size. Need to block /24s - /32s.

upvoted 1 times

ZamanR Most Recent 5 days, 22 hours ago

A is the correct answer

“Blocking all subnets of 172.16.0.0/16 major network with a mask lower than 23 from coming in”

would block 172.16.16.0/20.

The first prefix-list “ip prefix-list PL-1 deny 172.16.0.0/16 le 23” means “all networks that fall within

the 172.16.0.0/16 range AND that have a subnet mask of /23 or less” are denied.

The second prefix-list “ip prefix-list PL-1 permit 0.0.0.0/0 le 32” means allows all other prefixes.

upvoted 1 times

louisvuitton12 1 month, 3 weeks ago

Selected Answer: A

In summary, any subnet mask with a number higher than 23 (like /24, /25, /26, etc.)

upvoted 1 times

night_wolf_in 1 month, 3 weeks ago

Selected Answer: B

Block subnets smaller than 23, meaning 24,25, etc.
https://www.ciscozine.com/cisco-prefix-lists/

upvoted 1 times

BTK0311 3 months ago

The best configuration to block all subnets of the 172.16.0.0/16 major network with a mask lower than /23 from being advertised by R2 to R1 is option B:

B. ip prefix-list PL-1 deny 172.16.0.0/16 le 23 ip prefix-list PL-1 permit 0.0.0.0/0 le 32 ! router bgp 100 neighbor 192.168.100.2 remote-as 200 neighbor 192.168.100.2 prefix-list PL-1 in

This configuration uses a prefix-list (PL-1) to deny routes with a prefix length less than or equal to /23 from the 172.16.0.0/16 major network. It then permits all other routes. The prefix-list PL-1 is applied to the BGP neighbor 192.168.100.2 in the inbound direction using the prefix-list PL-1 in command.

Option A, C, and D either don't specify the correct prefix-list filtering criteria or use access-lists, which are not the most appropriate for this task. Option B aligns with the requirement to block subnets with a mask lower than /23 from the major network.

upvoted 1 times

JieW 3 months, 3 weeks ago

Selected Answer: A

Ge 23 Le 32 means 23-32. when it states lower than a subnet, it means lower number.
i encourage all to research what that means.
https://learningnetwork.cisco.com/s/question/0D53i00000Kt3t5CAB/ge-le

upvoted 1 times

chris110 3 months, 3 weeks ago

Selected Answer: B

To block all subnets of 172.16.0.0/16 with a mask lower than 23 from coming in on R1, you can use either a prefix-list or an access list. Let's evaluate the provided options:

A. This option uses a prefix-list and denies subnets of 172.16.0.0/16 with a mask greater than or equal to 23. This is incorrect because you want to block subnets with a mask lower than 23.

B. This option uses a prefix-list and denies subnets of 172.16.0.0/16 with a mask less than or equal to 23. This is the correct option because it matches the requirement.

C. This option uses a prefix-list but doesn't specify the mask length in the deny statement, so it would not block any specific subnets within 172.16.0.0/16.

D. This option uses an access list but denies subnets of 172.16.0.0/16 with a mask of 0.0.254.255, which is not the correct mask to block subnets with a mask lower than 23.

So, the correct configuration is option B

upvoted 2 times

chris110 3 months, 2 weeks ago

```
ip prefix-list PL-1 deny 172.16.0.0/16 le 23
ip prefix-list PL-1 permit 0.0.0.0/0 le 32

router bgp 100
neighbor 192.168.100.2 remote-as 200
neighbor 192.168.100.2 prefix-list PL-1 in
```

This configuration will block all subnets of 172.16.0.0/16 with a mask lower than /23 from being advertised from R2 to R1.

upvoted 1 times

siyamak 4 months, 1 week ago

The answer is B

upvoted 1 times

HarwinderSekhon 4 months, 1 week ago

CCNP is more of English exam vs networking :P

upvoted 4 times

Commando1664 5 months ago

How can it be A when it says pemit 172.16.0.0/16 with a subnet mask greater than or equal to 23...It's B.

upvoted 3 times

inteldarvid 5 months ago

Selected Answer: A

Sorry, i understand the option correct is A

upvoted 1 times

[Removed] 5 months, 1 week ago

Selected Answer: A

As said in other comments. This seems to be an english trick question. But like Dacusai explained:

When talking about a network that is lower (smaller) than /23, then you have to think of prefix /24 through /32, these broadcast domains are smaller than the broadcast domain of a /23 prefix.
If we block network 172.16.0.0/16 prefix less than or equal to /23, then we are blocking /22, /21, /20, etc, up to /16 and permitting everything else.
These network are very large network and we are left with what would be a large RIB of /24 networks. This does not improve convergence.

I also got tricked into it and initially answered B.

upvoted 3 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

the answer correct is B, because is lower than /23 is similar (le 23). Option B

upvoted 1 times

sajjad_gayyem 5 months, 3 weeks ago

Selected Answer: B

I guess lower than 23 mean 22 and 21.

upvoted 1 times

Almylle 5 months, 3 weeks ago

Im not sure, lower subnets mean 23, 24, 25 ,26 and so on, so in this case means ge in the prefix list

upvoted 1 times

Dacusai 8 months ago

LE /23 it means that all subnets including /16 will be deny.
GE /23 it means that from there to /32 will be deny, so the one with /26 will be permitted,
LE/32 will permit the rest of the traffic

upvoted 1 times

HungarianDish 8 months ago

I guess, the question should state "with a mask of 23 or lower". "ge /23" means euqal or greater than /23.
172.16.0.0/16 ge 23
= match any IPs from range 172.16.0.0 - 172.16.255.255 as long as subnet is /23 or greater
= match any IPs from 172.16.0.0/16 with mask euqal or greater than /23
= match 172.16.0.0/16 network and all prefixes contained therein with a length of 23 or greater
= match all prefixes within the 172.16.0.0/16 network that are at least 23 bits in length

upvoted 1 times

Question #81 *Topic 1*

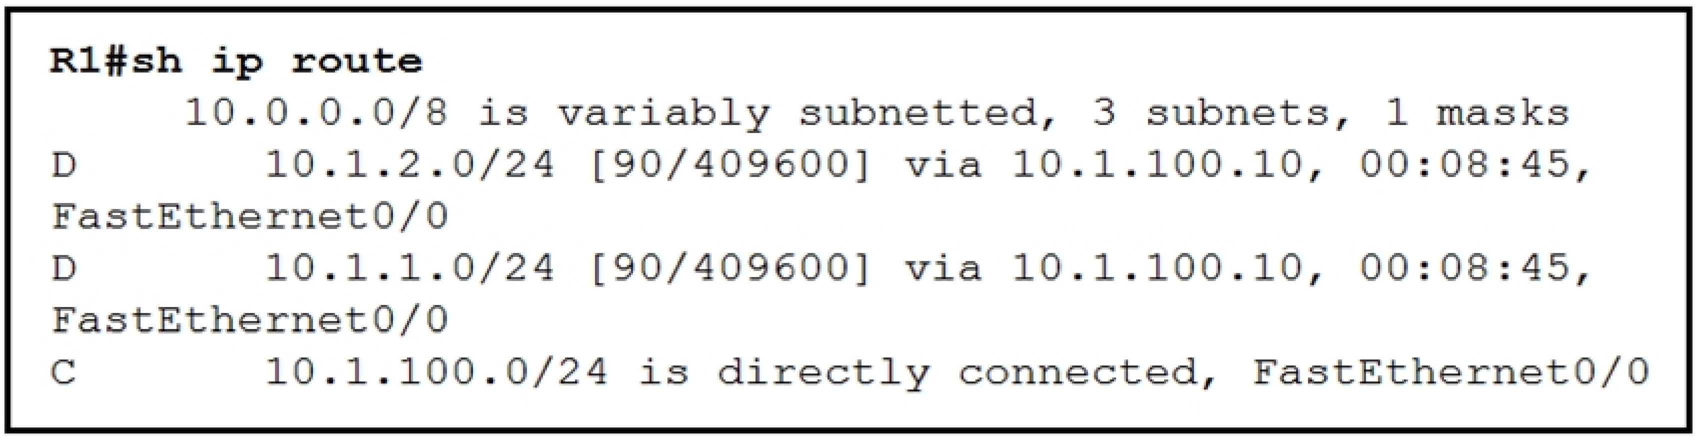

```
R1#sh ip route
     10.0.0.0/8 is variably subnetted, 3 subnets, 1 masks
D       10.1.2.0/24 [90/409600] via 10.1.100.10, 00:08:45,
FastEthernet0/0
D       10.1.1.0/24 [90/409600] via 10.1.100.10, 00:08:45,
FastEthernet0/0
C       10.1.100.0/24 is directly connected, FastEthernet0/0
```

Refer to the exhibit. An engineer configures the router 10.1.100.10 for EIGRP autosummarization so that R1 should receive the summary route of 10.0.0.0/8.
However, R1 receives more specific /24 routes.
Which action resolves this issue?

A. Router R1 should configure ip summary address eigrp (AS number) 10.0.0.0 255.0.0.0 for the R1 Fast Ethernet 0/0 connected interface.

B. Router R1 should configure ip route 10.0.0.0 255.0.0.0 null 0 for the routes that are received on R1.

C. Router 10.1.100.10 should configure ip route 10.0.0.0 255.0.0.0 null 0 for the routes that are summarized toward R1.

D. Router 10.1.100.10 should configure ip summary address eigrp (AS number) 10.0.0.0 255.0.0.0 for the R1 Fast Ethernet 0/0 connected interface.

Correct Answer: *D*

Community vote distribution

D (100%)

Huntkey Highly Voted 1 year, 3 months ago
EIGRP with auto-summary turned on, will not auto-summarize external routes, or routes learned from another router. It only summarizes locally injected internal routes. The summary is only generated at the major network boundary. For example, if there is 10.10.10.0/24 locally injected (with network statement) in the EIGRP, but the transit link to the EIGRP neighbor is in 10.0.0.0/24, then the 10.0.0.0/8 summary is not generated.
upvoted 6 times

jansan55 Most Recent 3 months, 2 weeks ago
Selected Answer: D
A good explanation:
https://study-ccna.com/eigrp-automatic-manual-summarization/
upvoted 2 times

inteldarvid 5 months, 2 weeks ago
Selected Answer: D
D is correct
upvoted 2 times

Question #82

Topic 1

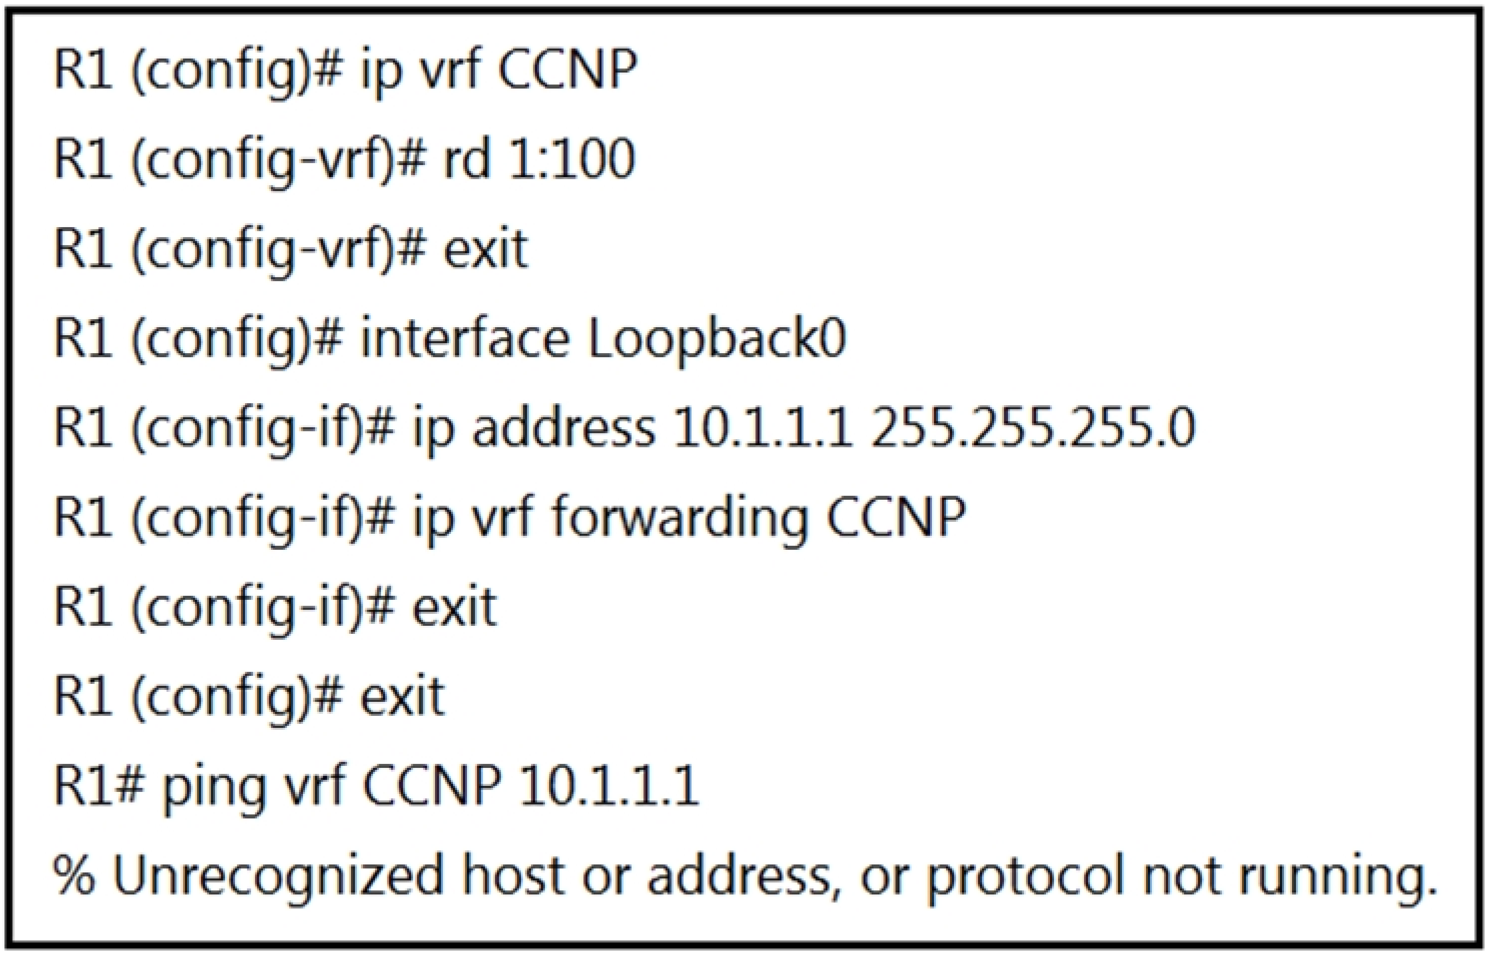

```
R1 (config)# ip vrf CCNP
R1 (config-vrf)# rd 1:100
R1 (config-vrf)# exit
R1 (config)# interface Loopback0
R1 (config-if)# ip address 10.1.1.1 255.255.255.0
R1 (config-if)# ip vrf forwarding CCNP
R1 (config-if)# exit
R1 (config)# exit
R1# ping vrf CCNP 10.1.1.1
% Unrecognized host or address, or protocol not running.
```

Refer to the exhibit. Which command must be configured to make VRF CCNP work?

A. interface Loopback0 ip address 10.1.1.1 255.255.255.0 vrf forwarding CCNP

B. interface Loopback0 ip address 10.1.1.1 255.255.255.0

C. interface Loopback0 vrf forwarding CCNP

D. interface Loopback0 ip address 10.1.1.1 255.255.255.0 ip vrf forwarding CCNP

Correct Answer: *B*

Reference:

https://community.cisco.com/t5/mpls/interface-ip-removed-after-apply-the-ip-vrf-forwarding/td-p/487122

Community vote distribution

B (85%) C (15%)

ZamanR 5 days, 5 hours ago

B is the correct answer

upvoted 1 times

Normanby 3 weeks, 4 days ago

Selected Answer: B

It is because they typed the VRF command AFTER the ip address, the ip got removed, so we need to add it back...

upvoted 1 times

Brand 3 months, 3 weeks ago

Selected Answer: B

Just put the IP back to the interface and you're good.

```
R1(config)#ip vrf CCNP
R1(config-vrf)#rd 1:100
R1(config-vrf)#exit
R1(config)#int lo 0
R1(config-if)#
R1(config-if)#ip address 10.1.1.1 255.255.255.0
R1(config-if)#ip vrf forwarding CCNP
% Interface Loopback0 IPv4 disabled and address(es) removed due to disabling VRF CCNP
R1(config-if)#exit
R1(config)#do show run int lo 0
Building configuration...

Current configuration : 66 bytes
!
interface Loopback0
ip vrf forwarding CCNP
```

no ip address
end

upvoted 4 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

correct B

upvoted 3 times

sajjad_gayyem 5 months, 3 weeks ago

Selected Answer: B

like Dacusai comment.

upvoted 2 times

Dacusai 8 months ago

Just another cisco thing, when you add the vrf command on an int the ip is removed so you need to added after, the correct answer should be the vrf command follow by the Ip int command. Tricky.

upvoted 4 times

Xerath 10 months ago

Selected Answer: B

After adding the int to the VRF, the assigned IP address is removed, so we just need to reconfigure the IP address on that same int.

upvoted 3 times

Lilienen 10 months, 2 weeks ago

Selected Answer: B

B is correct, as the vrf statement removes the IP address from the interface, therefore we have to re-add the IP address command

upvoted 4 times

PimplePooper 11 months, 2 weeks ago

Selected Answer: C

This is a stupid question. The questions asks for the VRF CCNP to work and not for the loopback ip address to work by its own. Answer B - if you want just the loopback to work without a vrf and answer C - for the loopback interface to be part of a vrf.

upvoted 3 times

Noproblem22 1 year, 1 month ago

D is the correct answer. Once you assign the ip address to lookback interface, you need to enable the "ip vrf forwarding CCNP".

upvoted 1 times

> **[Removed]** 1 year ago
>
> I believe adding “ip vrf forwarding CCNP” to an interface will remove the current IP address. Therefore the IP address statement must be added afterward. In the example, the vrf forwarding CCNP already now exists on the interfaces, but it lacks an IP address. This leaves us with needed to only re-add the IP address that is desired.
>
> upvoted 6 times

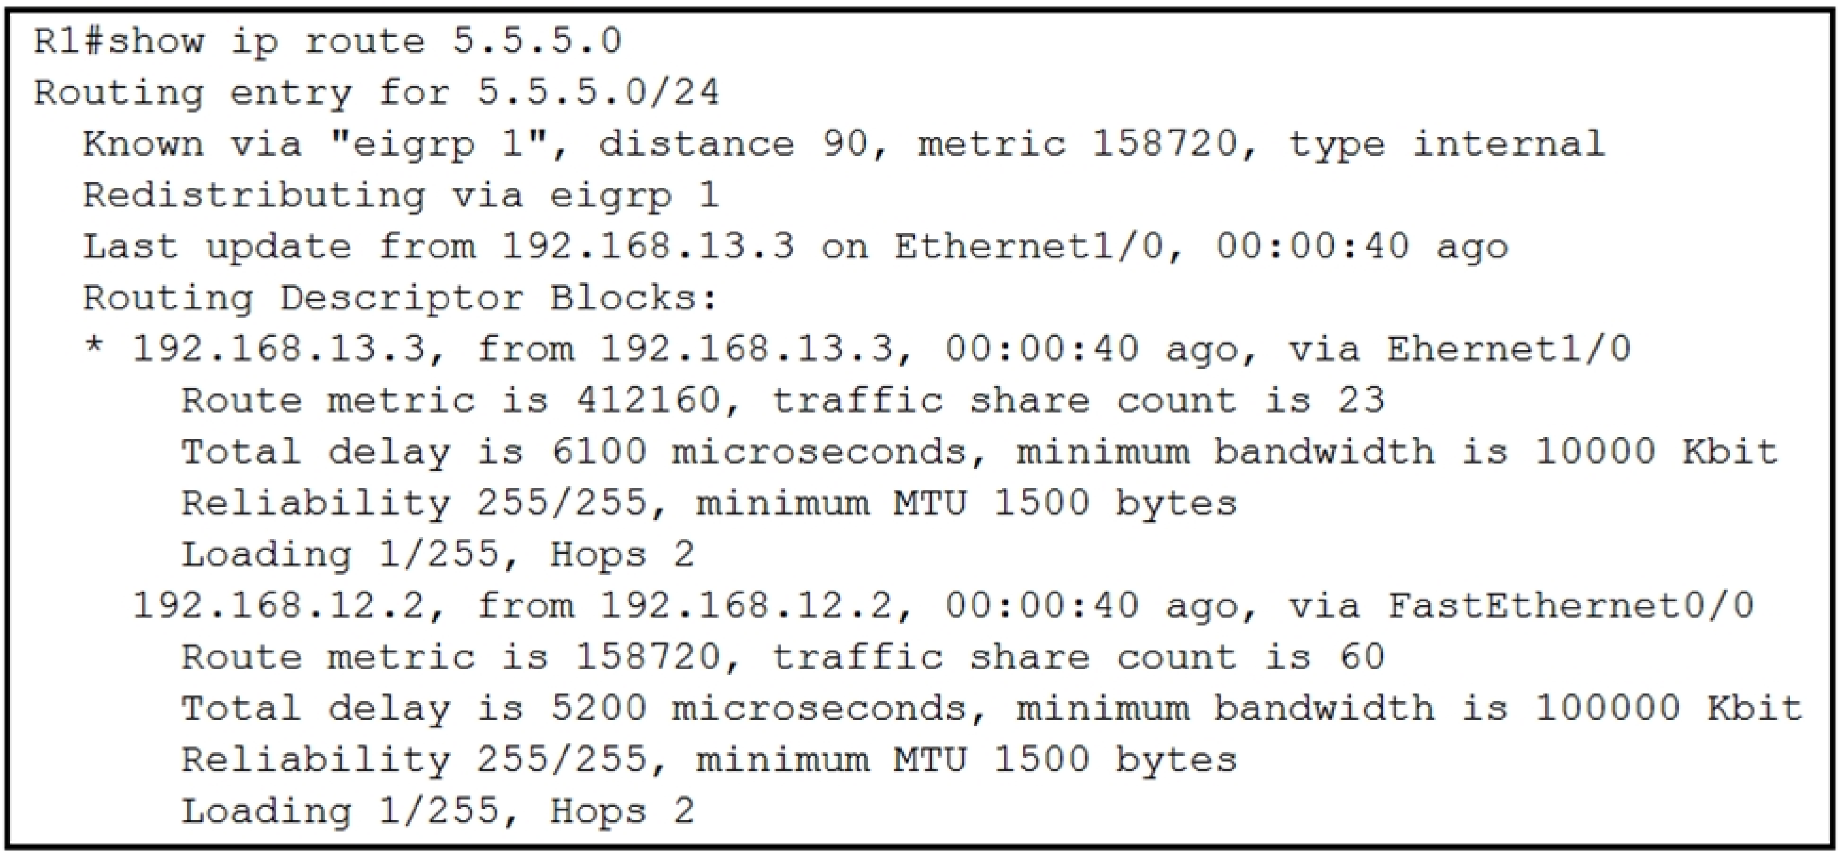

```
R1#show ip route 5.5.5.0
Routing entry for 5.5.5.0/24
  Known via "eigrp 1", distance 90, metric 158720, type internal
  Redistributing via eigrp 1
  Last update from 192.168.13.3 on Ethernet1/0, 00:00:40 ago
  Routing Descriptor Blocks:
  * 192.168.13.3, from 192.168.13.3, 00:00:40 ago, via Ehernet1/0
      Route metric is 412160, traffic share count is 23
      Total delay is 6100 microseconds, minimum bandwidth is 10000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
    192.168.12.2, from 192.168.12.2, 00:00:40 ago, via FastEthernet0/0
      Route metric is 158720, traffic share count is 60
      Total delay is 5200 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
```

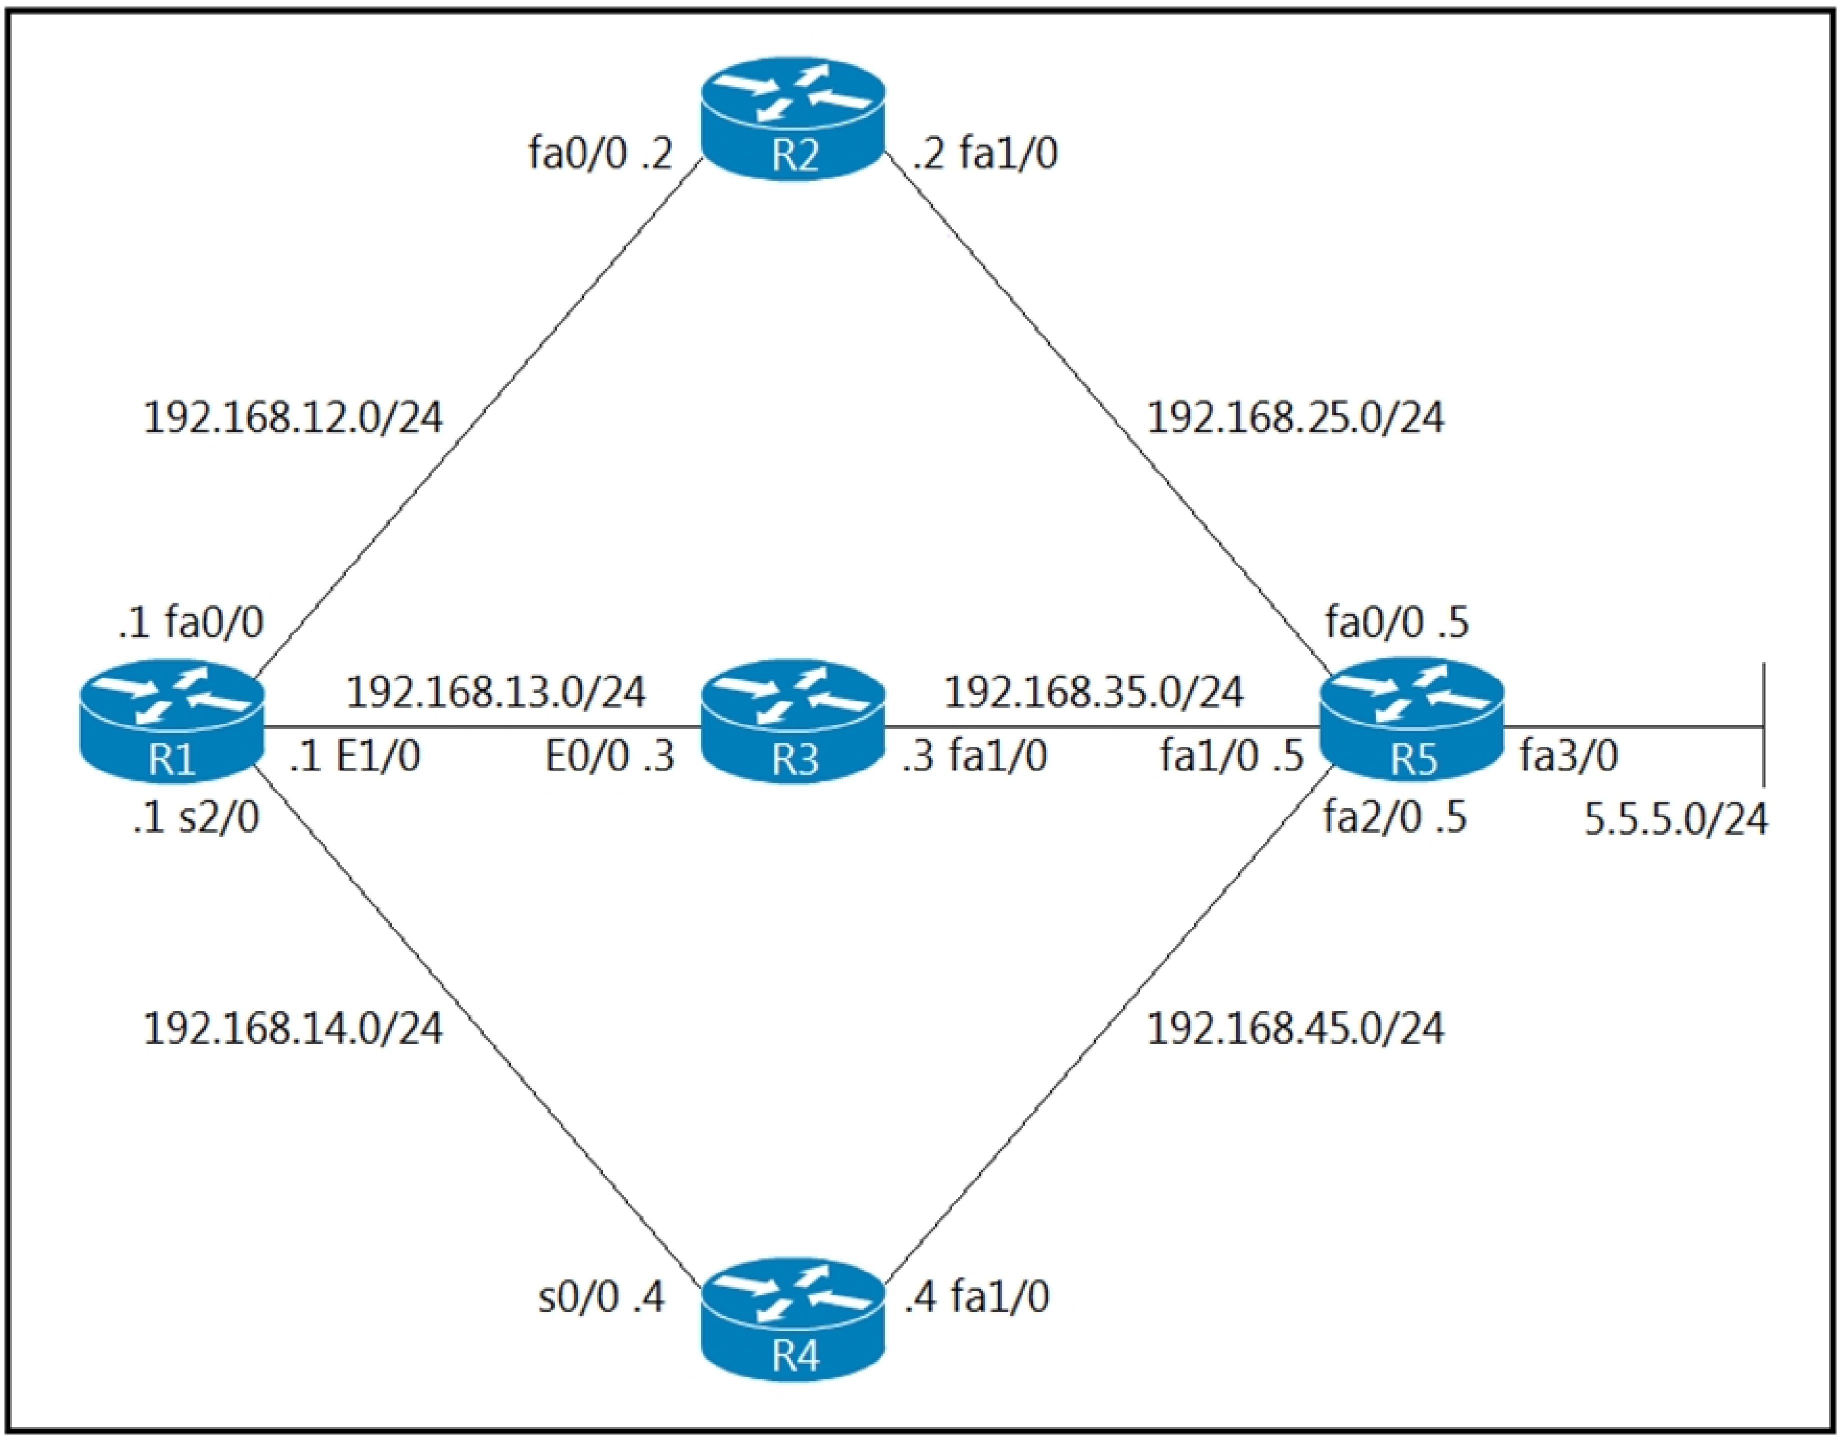

Refer to the exhibits. An engineer investigates a routing issue on R1 and finds that traffic destined to 5.5.5.0/24 does not take all of the paths. Which action resolves the issue?

A. Increase the variance value in EIGRP.

B. Decrease the variance value in EIGRP.

C. Remove the adjacency of R3 from EIGRP.

D. Stop advertising 192.168.13.0/24 in EIGRP.

Correct Answer: *A*

Reference:

https://community.cisco.com/t5/networking-documents/troubleshooting-eigrp-variance-command/ta-p/3129662#:~:text=EIGRP%20provides%20a%20mechanism%20to,means%20equal%2Dcost%20load%20balancing

HungarianDish 6 months, 4 weeks ago

FD of feasible successor / FD of successor ≈ variance
412160 / 158720 = 2.59 -> variance = 3

upvoted 1 times

Dataset 1 year, 4 months ago

Correcto A
EIGRP variante enables unequal cost path balance routing and add those prefixes to the EIGRP routing table.
Regards

upvoted 1 times

Question #84

Topic 1

DRAG DROP -

Drag and drop the MPLS VPN concepts from the left onto the correct descriptions on the right.

Select and Place:

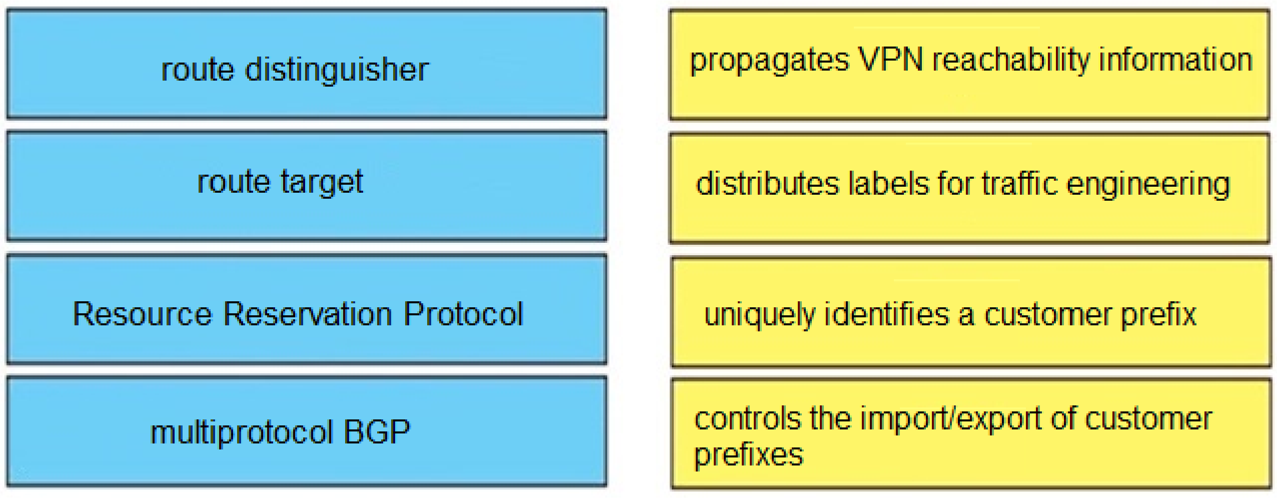

route distinguisher	propagates VPN reachability information
route target	distributes labels for traffic engineering
Resource Reservation Protocol	uniquely identifies a customer prefix
multiprotocol BGP	controls the import/export of customer prefixes

Correct Answer:

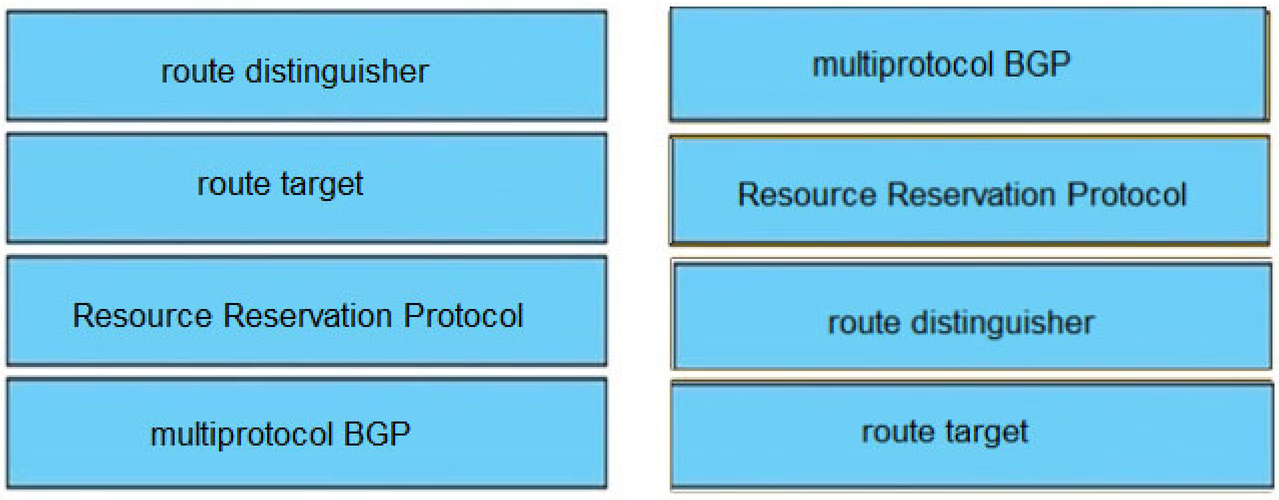

route distinguisher	multiprotocol BGP
route target	Resource Reservation Protocol
Resource Reservation Protocol	route distinguisher
multiprotocol BGP	route target

⊟ **Dave513** Highly Voted 3 years ago

Route distinguisher - Uniquely identifies a customer prefix;
Route Target - Controls the import/export of customer prefixes;
Resource Reservation Protocol - Distributes labels for traffic engineering;
Multi-Protocol BGP - Propagates VPN reachability information.

upvoted 9 times

⊟ **error_909** Most Recent 2 years, 3 months ago

The given answer is correct

upvoted 1 times

⊟ **examShark** 2 years, 4 months ago

The given answer is correct

upvoted 1 times

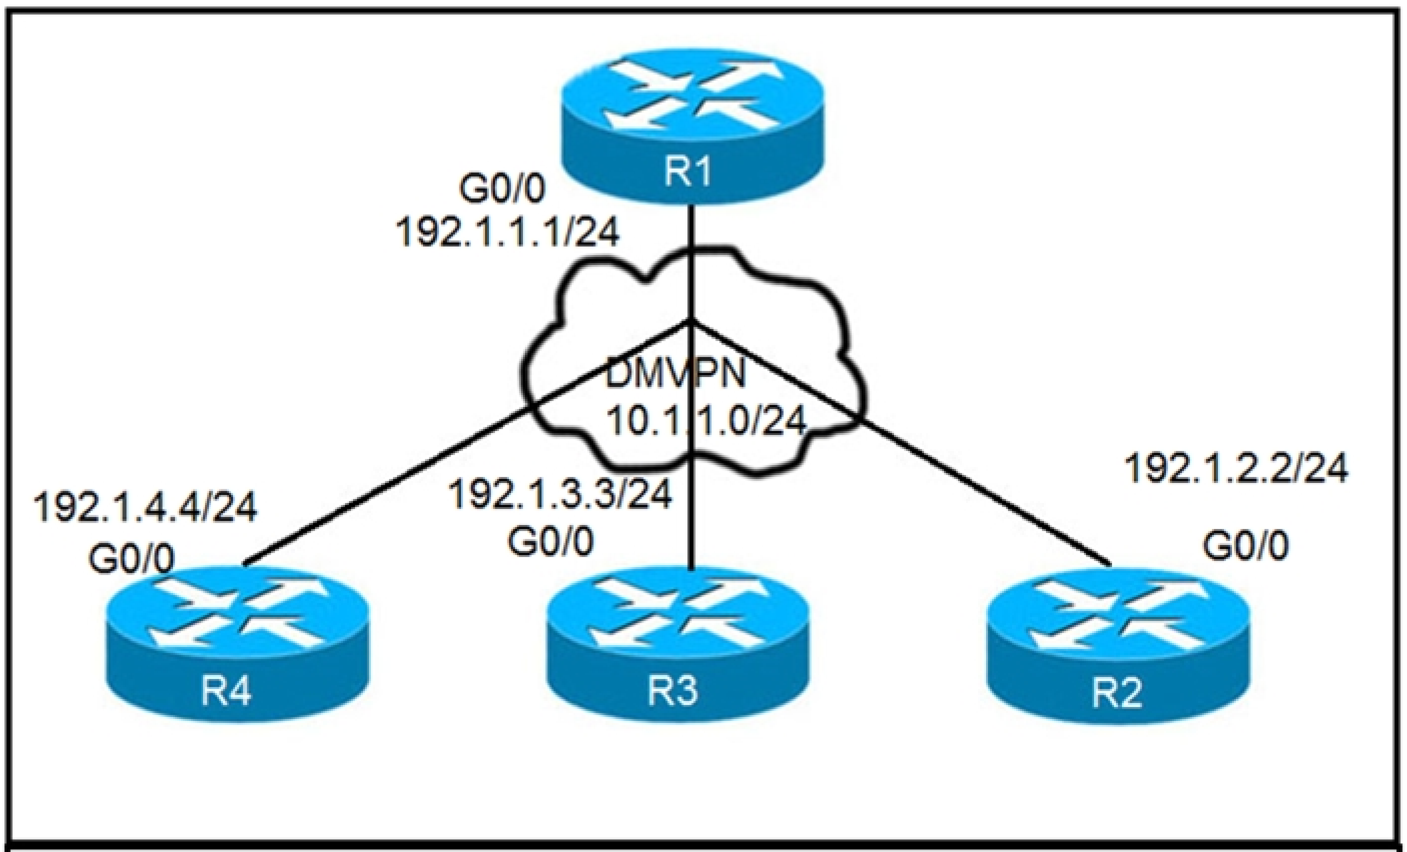

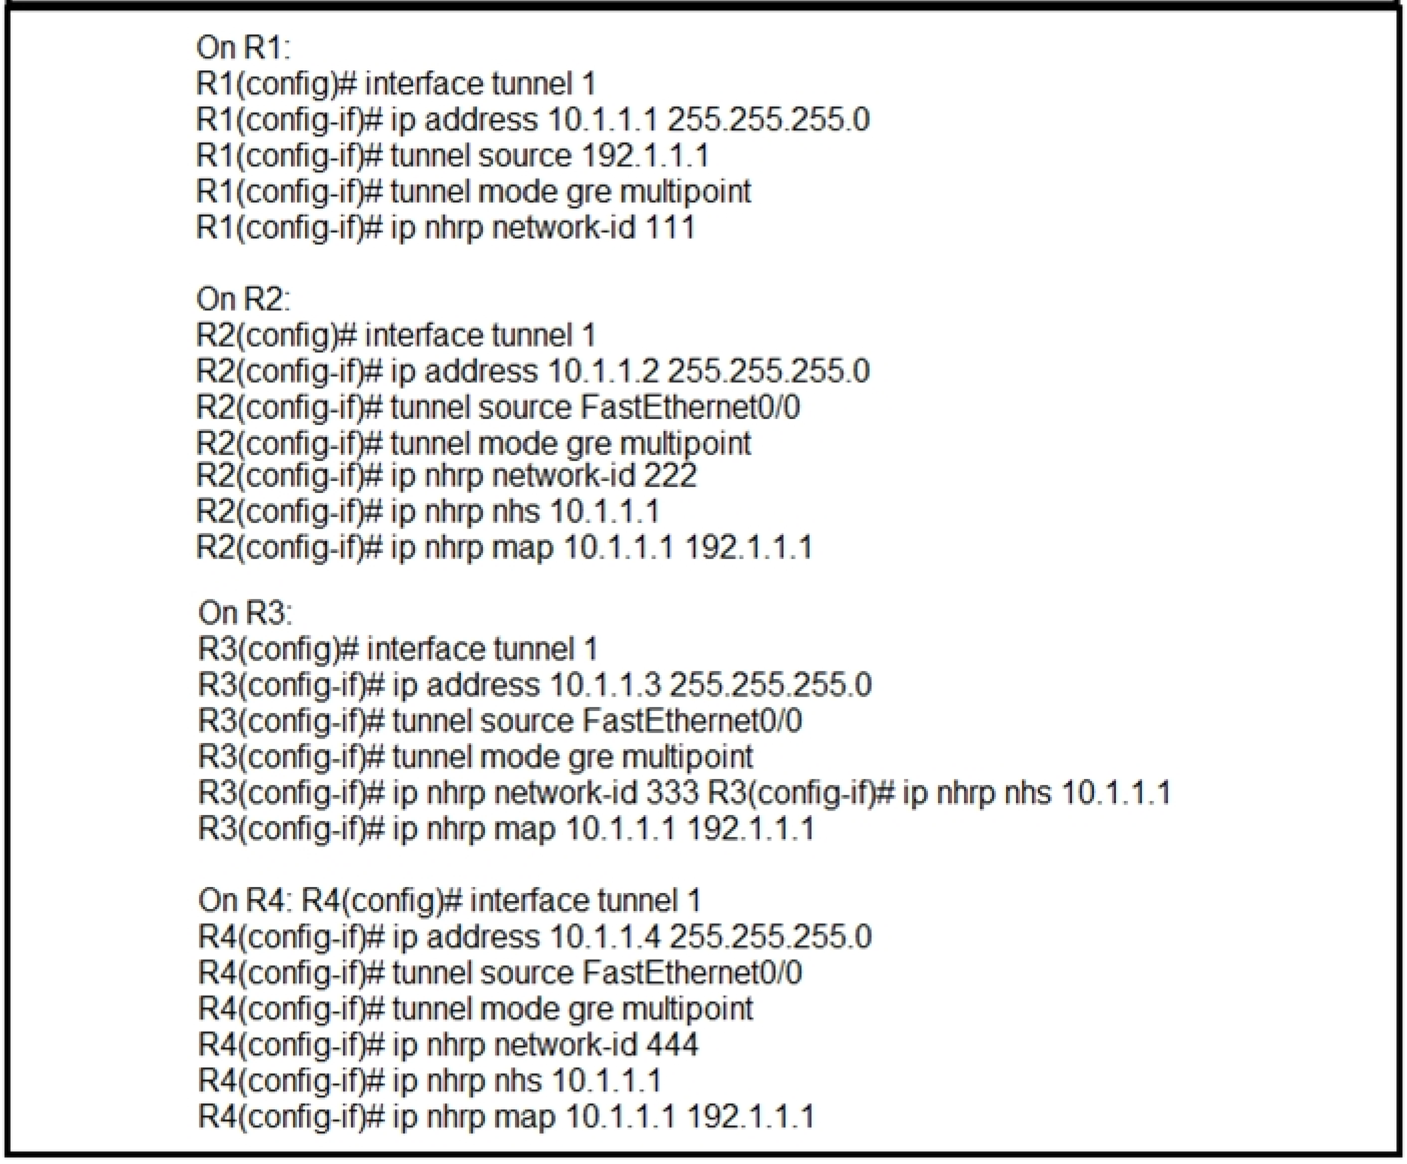

```
On R1:
R1(config)# interface tunnel 1
R1(config-if)# ip address 10.1.1.1 255.255.255.0
R1(config-if)# tunnel source 192.1.1.1
R1(config-if)# tunnel mode gre multipoint
R1(config-if)# ip nhrp network-id 111

On R2:
R2(config)# interface tunnel 1
R2(config-if)# ip address 10.1.1.2 255.255.255.0
R2(config-if)# tunnel source FastEthernet0/0
R2(config-if)# tunnel mode gre multipoint
R2(config-if)# ip nhrp network-id 222
R2(config-if)# ip nhrp nhs 10.1.1.1
R2(config-if)# ip nhrp map 10.1.1.1 192.1.1.1

On R3:
R3(config)# interface tunnel 1
R3(config-if)# ip address 10.1.1.3 255.255.255.0
R3(config-if)# tunnel source FastEthernet0/0
R3(config-if)# tunnel mode gre multipoint
R3(config-if)# ip nhrp network-id 333 R3(config-if)# ip nhrp nhs 10.1.1.1
R3(config-if)# ip nhrp map 10.1.1.1 192.1.1.1

On R4: R4(config)# interface tunnel 1
R4(config-if)# ip address 10.1.1.4 255.255.255.0
R4(config-if)# tunnel source FastEthernet0/0
R4(config-if)# tunnel mode gre multipoint
R4(config-if)# ip nhrp network-id 444
R4(config-if)# ip nhrp nhs 10.1.1.1
R4(config-if)# ip nhrp map 10.1.1.1 192.1.1.1
```

Refer to the exhibits. Phase-3 tunnels cannot be established between spoke-to-spoke in DMVPN. Which two commands are missing? (Choose two.)

A. The ip nhrp redirect command is missing on the spoke routers.

B. The ip nhrp shortcut command is missing on the spoke routers.

C. The ip nhrp redirect command is missing on the hub router.

D. The ip nhrp shortcut command is missing on the hub router.

E. The ip nhrp map command is missing on the hub router.

Correct Answer: *BC*

Community vote distribution

BC (75%) 8% Other

Malasxd Highly Voted 1 year, 10 months ago

Selected Answer: BC

Well, it's a really complicate question. For me, it's missing 3 commands. The "ip nghrp map multicast dynamic" and "ip nhrp redirect" in the hub. And the IP nhrp shortcut in the spokes. In a lab, the DMVPN Phase wouldn't work without theses 3 commands.
BC seems correct, but without E it wouldn't work. It's complicated.

upvoted 5 times

Colmenarez Most Recent 3 months, 3 weeks ago

Selected Answer: BC

OCG Pag 773

"The Phase 3 DMVPN configuration for the hub router adds the interface parameter command IP NHRP REDIRECT on the hub router. This command checks the flow of packets on the tunnel interface and sends a redirect message to the source spoke router when it detects packets hairpinning out of the DMVPN cloud. Hairpinning means that traffic is received and sent out an interface in the same cloud (identified by the NHRP network ID). For instance, hairpinning occurs when packets come in and go out the same tunnel interface.
The Phase 3 DMVPN configuration for spoke routers uses the mGRE tunnel interface and uses the command IP NHRP SHORTCUT on the tunnel interface.

upvoted 1 times

HungarianDish 8 months ago

Selected Answer: BC

Hub(config)#interface tunnel 1
Hub(config-if)#ip nhrp redirect
-> Hub notifies spoke routers of suboptimal traffic paths

Spoke1(config)#interface tunnel 1
Spoke1(config-if)#ip nhrp shortcut
-> Spokes send a resolution request for a shortcut path after receiving an NHRP redirect traffic indication message

upvoted 3 times

dq28 11 months, 3 weeks ago

Selected Answer: BD

Shortcut is needed on hub and spokes to initiate spoke-to-spoke tunnels. Redirect is an optimization that is needed for spoke-to-spoke communication only if you have summarized routes, if the spoke has the complete routing-table it is not needed.

https://www.cisco.com/c/en/us/support/docs/security/dynamic-multipoint-vpn-dmvpn/211292-Configure-Phase-3-Hierarchical-DMVPN-wit.html#anc9
https://www.ciscozine.com/dmvpn-phase-3-guide/

upvoted 1 times

chris7890 1 year ago

The given answer is correct. Phase 3: Use ip nhrp redirect on hub routers & ip nhrp shortcuts on spoke routers.
https://network-insight.net/2015/02/03/design-guide-dmvpn-phases/

upvoted 1 times

quyle 1 year, 2 months ago

I think question have a problem ip nhrp network-id different on int tun 1 cause dmvpn not up.

if question is true. I choose B and C. shortcut and redirect is representative of dmvpn phase 3

upvoted 1 times

M_Abdulkarim 1 year, 4 months ago

also ip nhrp map multicast dynamic is missing.
it allows NHRP to automatically add spoke routers to the multicast NHRP mappings.

upvoted 2 times

Hack4 1 year, 10 months ago

B&C are the best answers

upvoted 1 times

JingleJangus 1 year, 10 months ago

Selected Answer: CE

For phase 3 dmvpn you do. You absolutely do. Got cisco press right in front of me. Nhrp provides a mapping service of the tunnel ip to the nbma ip. So if you want those juicy tunnel to nbma mappings that the hub keeps in its pockets to know when a spoke could use a more optimal path (ip nhrp redirect) it will forward the redirect packet to the destination spoke and it will resolve with the remote spoke. Ip nhrp shortcut is enabled by default, so D is not correct.

upvoted 1 times

[Removed] 1 year, 10 months ago

Yea this is very misleading. For phase 3 spoke to spoke tunnels you have to have redirect on the hub and shortcut on the spokes. Yes the hub is missing the map command but this is specifically asking for phase 3 configurations.

upvoted 1 times

JingleJangus 1 year, 10 months ago

Selected Answer: DE

For phase 3 dmvpn you do. You absolutely do. Got cisco press right in front of me. Nhrp provides a mapping service of the tunnel ip to the nbma ip. So if you want those juicy tunnel to nbma mappings that the hub keeps in its pockets to know when a spoke could use a more optimal path (ip nhrp redirect) it will forward the redirect packet to the destination spoke and it will resolve with the remote spoke. Ip nhrp shortcut is enabled by default, so D is not correct.

upvoted 1 times

error_909 2 years, 3 months ago

B&C&D are all missing.
But since the question is asking specifically about Phase 3 we assume that
B&C are the write answers

upvoted 2 times

AliMo123 2 years, 4 months ago

wrong answers
E is correct since the Hub is missing command (ip NHRP map multicast dynamic) otherwise point-to point tunnel btw spokes wont form.

upvoted 1 times

Raider1 2 years, 2 months ago

You don't need NHRP map on the Hub

upvoted 2 times

JingleJangus 1 year, 10 months ago

For phase 3 dmvpn you do. You absolutely do. Got cisco press right in front of me. Nhrp provides a mapping service of the tunnel ip to the nbma ip. So if you want those juicy tunnel to nbma mappings that the hub keeps in its pockets to know when a spoke could use a more optimal path (ip nhrp redirect) it will forward the redirect packet to the destination spoke and it will resolve with the remote spoke. Ip nhrp shortcut is enabled by default, so D is not correct.

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct.
HUB: redirect, Spoke: shortcut for phase 3 dmvpn

upvoted 1 times

azharken 2 years, 7 months ago

ip nhrp map command is also required on hub router to establish basic dmvpn connection, redirect and shortcut commands comes afterwards

upvoted 1 times

Question #86 *Topic 1*

Which protocol is used to determine the NBMA address on the other end of a tunnel when mGRE is used?

A. NHRP

B. IPsec

C. MP-BGP

D. OSPF

Correct Answer: *A*

Community vote distribution

A (100%)

Xerath 10 months ago

Selected Answer: A

The given answer is correct

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #87

Topic 1

A DMVPN single hub topology is using IPsec + mGRE with OSPF.
What should be configured on the hub to ensure it will be the designated router?

A. route map to set the metrics of learned routes to 110

B. tunnel interface of the hub with ip nhrp ospf dr

C. OSPF priority to 0

D. OSPF priority greater than 1

Correct Answer: *D*

Community vote distribution

D (100%)

Normanby 3 weeks, 4 days ago

Selected Answer: D

Classic misdirect question - the question really is: how to force a router to be the DR.

upvoted 1 times

chris110 3 months, 3 weeks ago

In a DMVPN (Dynamic Multipoint Virtual Private Network) single hub topology with OSPF, to ensure that the hub router becomes the designated router (DR) for the OSPF network, you need to configure the OSPF priority of the hub router to be greater than 0. OSPF uses the priority value to determine which router becomes the DR. The router with the highest priority becomes the DR.

So, the correct option is:

D. OSPF priority greater than 1

Setting the OSPF priority to a value greater than 1 will increase the likelihood of the hub router becoming the designated router in the OSPF network. Typically, setting the priority to 1 means the router is not eligible to become the DR, so it should be set to a value greater than 1 to become the DR.

upvoted 2 times

[Removed] 2 years ago

A priority of 0 means a router will not participate in DR/BDR election.

upvoted 3 times

examShark 2 years, 4 months ago

The given answer is correct.
o means never DR, 1 is the default, highest wins

upvoted 4 times

Question #88 *Topic 1*

What are two purposes of using IPv4 and VPNv4 address-family configurations in a Layer 3 MPLS VPN? (Choose two.)

A. RD is prepended to the IPv4 route to make it unique.

B. The VPNv4 address consists of a 64-bit route distinguisher that is prepended to the IPv4 prefix.

C. MP-BGP is used to allow overlapping IPv4 addresses between customers to advertise through the network.

D. The IPv4 address is needed to tag the MPLS label.

E. The VPNv4 address is used to advertise the MPLS VPN label.

Correct Answer: *AB*

Community vote distribution

AE (55%) AB (36%) 5%

⊟ **HungarianDish** Highly Voted 6 months, 3 weeks ago

Selected Answer: AE

A:
-64-bit RD prepended to IPv4 prefix to make customer routes unique = VPNv4 address
https://community.cisco.com/t5/switching/i-am-not-clear-on-the-difference-between-ipv4-and-vpnv4-address/td-p/2463679
E:
The VPN label is advertised to all other PE routers in an MP-BGP update.
https://www.ccexpert.us/mpls/vpn-label-propagation.html

upvoted 8 times

⊟ **Pietjeplukgeluk** 1 month ago

E == The VPNv4 address is used to advertise the MPLS VPN label.
The VPNv4 address does NOT include ANY VPN label, also the VPNv4 address is just a name for an 48bit RD and a 32 bit prefix. The actual MP-BGP advertisement will just include the following:
1. RD (Route Distinguisher)
2. IPv4 prefix
3. Next Hop
4. VPN Label

Instead of E, i personally think B is the better answer. B is still partly wrong as B states the word "prepended"==WRONG the IPv4 prefix is actually "appended" NOT "prepended". Anyway, another question of bad quality.
ref: https://networklessons.com/cisco/ccnp-enarsi-300-410/mpls-layer-3-vpn-explained#RD_Route_Distinguisher

upvoted 2 times

⊟ **HungarianDish** 6 months, 2 weeks ago

https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/5eU6DfQV/TECMPL-3201.pdf
-vpnv4 AFI for PE to PE (label information)
-All vpnv4 routes get an assigned label
-vpnv4 routes are exchanged between vpnv4 peers (PEs)
https://community.cisco.com/t5/mpls/mpls-vpn-inner-label/td-p/2356956
For every prefix in every VRF routing table, we make a corresponding vpnv4 prefix (RD is added to the IPv4 prefix).
BGP on the egress PE then assigns a label to that vpnv4 prefix.
The PE router then pushes this entry into the LFIB, with that label as incoming label (and also the label operation and next-hop).
BGP advertises the vpnv4 prefix + MPLS label to all other PE routers.
https://www.rfc-editor.org/rfc/rfc3107
When BGP is used to distribute a particular route, it can be also be used to distribute a Multiprotocol Label Switching (MPLS) label which is mapped to that route.

upvoted 1 times

⊟ **Tedmus** Most Recent ⊙ 1 month ago

Selected Answer: AC

A. Is clear true because the RD is prepended.
B. On the first view it seems to be true, BUT the wording is tricky. A VPNv4 prefix consits of the 64-bit RD and the 32-bit IPv4 prefix making it a 96-bit prefix.
C. It is true.
D. non-sense
E. non-sense

upvoted 1 times

⊟ **Tedmus** 1 month ago

Sry my fault. Ingore it.

upvoted 1 times

⊟ **louisvuitton12** 1 month, 3 weeks ago

Selected Answer: AE

Option B can NOT be the answer because RD is 48 bit.
In summary route-distinguisher: Specifies an RD, a string of 3 to 21 characters. An RD can be in one of the following formats:
16-bit AS number:32-bit user-defined number. For example, 101:3.
32-bit IP address:16-bit user-defined number. For example, 192.168.122.15:1.
32-bit AS number:16-bit user-defined number, where the AS number must be equal to or greater than 65536. For example, 65536:1.

upvoted 2 times

Pietjeplukgeluk 1 month ago

You indicate the RD==48 bits. The RD is clearly 64 bits, making B the correct answer. See https://networklessons.com/cisco/ccnp-enarsi-300-410/mpls-layer-3-vpn-explained#RD_Route_Distinguisher

upvoted 1 times

louisvuitton12 1 month, 3 weeks ago

PLEASE IGNORE THIS COMMENT, AB are correct my mistake.

upvoted 1 times

jansan55 3 months, 1 week ago

Selected Answer: BE

Something has changed in the order of answers?
For example HungarianDish explained why BE the right choice, while voted AE.
B: The VPNv4 address consists of a 64-bit route distinguisher that is prepended to the IPv4 prefix
E: The VPNv4 address is used to advertise the MPLS VPN label.

upvoted 1 times

jansan55 3 months, 1 week ago

Answer A wrong, because RD is prepended to the IPv4 prefix, not to the IPv4 route.

upvoted 1 times

Colmenarez 3 months, 3 weeks ago

Selected Answer: AB

OCG pag. 741-742

MPLS Layer 3 VPNv4 Address

Let's now go back to overlapping IPv4 address spaces. If all customer routes are being redistributed into MP-BGP, how does BGP handle identical network prefixes that belong to different customers? It uses a route distinguisher (RD) to expand the customer's IP prefix so that it includes a unique value that distinguishes it from the other identical prefixes. The RD is generated and used by the PE routers on a per-customer VRF instance basis, and to keep things simple, the RD is used regardless of whether there are overlapping address spaces. So, the RD is used all the time.

The unique 64-bit RD is prepended to the 32-bit customer prefix (IPv4 route) to create a 96-bit unique prefix called a VPNv4 address, as shown in Figure 18-14. This VPNv4 address is exchanged by the MP-IBGP neighboring routers.

upvoted 2 times

guy276465281819372 4 months, 1 week ago

Selected Answer: AB

A and B.

upvoted 1 times

JieW 4 months, 2 weeks ago

Selected Answer: AB

VPNv4 is known as an RD. They are not used to advertise the MPLS VPN Label

upvoted 2 times

inteldarvid 4 months, 2 weeks ago

Selected Answer: AB

A n B the give answer is correct, plese check anwser Rob_CCNP000 . I have the book. Its true

upvoted 1 times

adudeguy 6 months, 1 week ago

AB
E is wrong because MP-BGP (not VPNv4) is used to advertise MPLS VPN labels

upvoted 1 times

Rob_CCNP000 6 months, 1 week ago

Selected Answer: AB

Page 742 300-410 Official Cert Guide; RD (8-bytes) & IPv4 Address (4-Bytes) creates a 96-bit unique prefix called VPNv4 address.

upvoted 2 times

DenskyDen 6 months, 2 weeks ago

Selected Answer: AE

Read HungarianDish posted links.

upvoted 2 times

Dacusai 8 months ago

2 answers saying the same thing, no make sense.

upvoted 1 times

yeyuno 8 months, 2 weeks ago

https://community.cisco.com/t5/switching/i-am-not-clear-on-the-difference-between-ipv4-and-vpnv4-address/td-p/2463679

The VPNv4 address family is used to advertise VPNv4 NLRI. VPNv4 address consists of 64-bit Route Distinguisher (RD) prepended to IPv4 prefix. This is to make routes unique that are in different VRFs.

upvoted 1 times

GReddy2323 7 months, 3 weeks ago

So the answers are A & E?

upvoted 1 times

MarvinY 1 year, 12 months ago

Why C is wrong?

upvoted 3 times

gndrx78 1 year, 11 months ago

I think because it is VRF not MPLS allowing overlapping

upvoted 2 times

JingleJangus 1 year, 10 months ago

Yeah, like in the context of what we know mpls l3vpn to be, C is right, but its not in the context of the question. I had to re read C a couple times to see that.

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #89 *Topic 1*

What are two functions of MPLS Layer 3 VPNs? (Choose two.)

A. It is used for transparent point-to-multipoint connectivity between Ethernet links/sites.

B. A packet with node segment ID is forwarded along with shortest path to destination.

C. Customer traffic is encapsulated in a VPN label when it is forwarded in MPLS network.

D. BGP is used for signaling customer VPNv4 routes between PE nodes.

E. LDP and BGP can be used for Pseudowire signaling.

Correct Answer: *CD*

TheBaja 1 year, 1 month ago

C is OK.
D - it should be MBGP not BGP.

upvoted 2 times

Noproblem22 1 year, 1 month ago

The given answer looks good.

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 3 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #90 *Topic 1*

What are two MPLS label characteristics? (Choose two.)

A. The label edge router swaps labels on the received packets.

B. Labels are imposed in packets after the Layer 3 header.

C. LDP uses TCP for reliable delivery of information.

D. An MPLS label is a short identifier that identifies a forwarding equivalence class.

E. A maximum of two labels can be imposed on an MPLS packet.

Correct Answer: *CD*

Community vote distribution

CD (100%)

examShark Highly Voted 2 years, 4 months ago

The given answer is correct

upvoted 8 times

Ll123123 Most Recent 1 month, 4 weeks ago

A. The label edge router swaps labels on the received packets.
LSR swap label and LER take away the label or add the label when receive packet... but compare to other it seem this is more correct than other
B. Labels are imposed in packets after the Layer 3 header.
Label is layer 2.5
C. LDP uses TCP for reliable delivery of information.
LDP is udp
D. An MPLS label is a short identifier that identifies a forwarding equivalence class.
Mpls label consists of the label and a fec for traffic engineering use. Not only fec
E. A maximum of two labels can be imposed on an MPLS packet.
One mpls label for lfib routing, one mpls vpn label for ibgp routing

upvoted 1 times

mitosenoriko 11 months, 3 weeks ago

LDP discover use UDP after flow use TCP.
In this case decide TCP? I'm not sure.

upvoted 1 times

ahmeeedoox 1 year, 9 months ago

the question here is about Label !!!! so i would say the right answer are B and D
B because i think label came after the layer 3 header is added to the packet.
C it is talking about LDP not label !!!???

upvoted 1 times

larn 1 year, 7 months ago

MPLS Label is after the Layer2 Stack and Before the Layer 3 Stack

upvoted 4 times

steiger 2 years ago

Selected Answer: CD

voted for C and D

upvoted 2 times

error_909 2 years, 3 months ago

The given answer is correct C & D

upvoted 2 times

abc2k7 2 years, 5 months ago

A,D is true.

upvoted 1 times

gndrx78 1 year, 11 months ago

Label Edge routers add/remove labels, they do not swap

upvoted 4 times

vdsdrs 2 years, 4 months ago

No, Label SWITCH Router does label swaps

upvoted 1 times

Question #91 *Topic 1*

Which command allows traffic to load-balance in an MPLS Layer 3 VPN configuration?

A. multi-paths eibgp 2

B. maximum-paths 2

C. maximum-paths ibgp 2

D. multi-paths 2

Correct Answer: *C*

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/mpls/configuration/guide/mpls_cg/mp_vpn_multipath.html

Community vote distribution

C (60%) B (20%) 13% 7%

⊟ **ZachTL11** Highly Voted 2 years, 8 months ago

C is correct.
Remember MPLS Layer 3 VPN is used in an iBGP setup and not eBGP. This rules out B as the correct answer.

upvoted 13 times

⊟ **wts** 1 year, 3 months ago

MPLS L3VPN is used IGP/EBGP(CE-PE) and IGP+MPLS+MBGP[iBGP](PE-PE).
That is why the votes are divided.

upvoted 1 times

⊟ **louisvuitton12** Most Recent 1 month, 3 weeks ago

Selected Answer: C

Answer is C.

https://www.cisco.com/c/en/us/td/docs/dcn/nx-os/nexus9000/102x/configuration/label-switching/b-cisco-nexus-9000-series-nx-os-label-switching-configuration-guide-102x/m-configuring-mpls-layer-3-vpn-load-balancing.pdf

Example: MPLS Layer 3 VPN Load Balancing
The following example shows how to configure iBGP load balancing:
configure terminal
feature-set mpls
feature mpls l3vpn
feature bgp
router bgp 1.1
bestpath cost-community ignore
address-family ipv6 unicast
maximum-paths ibgp 4

upvoted 2 times

⊟ **BTK0311** 3 months ago

eBGP and iBGP Multipath Load Sharing Configuration Example
This following configuration example configures a router in IPv4 address-family mode to select two
BGP routes (eBGP or iBGP) as multipaths:
Device router bgp 40000
Deviceaddress-family ipv4 vrf RED
Devicemaximum-paths eibgp 2
Deviceend
This following configuration example configures a router in IPv6 address-family mode to select two
BGP routes (eBGP or iBGP) as multipaths:
Device router bgp 40000
Deviceaddress-family ipv6 vrf RED
Devicemaximum-paths eibgp 2

upvoted 1 times

⊟ **inteldarvid** 5 months, 2 weeks ago

Selected Answer: C

option corret is: C

https://www.cisco.com/c/en/us/td/docs/ios/12_2sx/feature/guide/fsxeibmp.html

The maximum-paths eibgp command used to configure Border Gateway Protocol (BGP) multipath load sharing in an Multiprotocol Label Switching (MPLS) virtual private network (VPN) using eBGP and iBGP routes. This feature is configured under a VPN routing and forwarding instance (VRF) in address family configuration mode. The number of multipaths is configured separately for each VRF. The number of paths that can be configured is determined by the version of Cisco IOS software

upvoted 1 times

Rob_CCNP000 6 months, 1 week ago

Selected Answer: C

The command “maximum-paths [ibgp] number-of-paths” configures the maximum number of multipaths allowed - MPLS uses ebgp so C is correct.

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

Selected Answer: C

https://www.cisco.com/c/en/us/td/docs/ios/12_2sx/feature/guide/fsxeibmp.html#wp1037690
This is how I see this based on the mentioned article. In BGP we can set different options for load sharing, depending on the design. With MPLS, we can configure "maximum-path ibgp" or "maximum-path eibgp", where the feature [ibgp] performs multipath forwarding only in PE-PE ibgp MPLS domain, and the feature [eibgp] can use both PE-PE ibgp and PE-CE ebgp domains for multipath calculations. If "maximum-paths eibgp 2" is not offered then "maximum-paths ibgp" is OK for MPLS (PE-PE) scenarios. = Answer "C"

upvoted 2 times

dancott 7 months ago

Selected Answer: C

Copied from Cisco config guide:
Example: MPLS Layer 3 VPN Load Balancing
The following example shows how to configure iBGP load balancing:

```
configure terminal
feature-set mpls
feature mpls l3vpn
feature bgp
router bgp 1.1
bestpath cost-community ignore
address-family ipv6 unicast
maximum-paths ibgp 4
```

upvoted 1 times

HungarianDish 8 months ago

Why aren't they offering this answer?

```
#address-family ipv4 vrf ...
#maximum-paths eibgp 2
```

-if the parameter "ibgp" is optional than answer "B" might be closer
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/mpls/configuration/guide/mpls_cg/mp_vpn_multipath.html
#maximum-paths [ibgp] "number-of-paths"

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

B) #maximum-paths 2 is only for PE-CE eBGP connections outside of the MPLS domain.

upvoted 1 times

Commando1664 8 months, 3 weeks ago

Summmary Steps from the article:
Configuring BGP Load Balancing for eBGP and iBGP

You can configure a Layer 3 VPN load balancing for an eBGP or iBGP network.
Prerequisites

Ensure that you are in the correct VDC (or use the switchto vdc command).
SUMMARY STEPS

1. configure terminal

2. feature- s et mpls

3. feature mpls l3vpn

4. feature bgp

5. router bgp as-number

6. (Optional) bestpath cost-community ignore

7. address-family { ipv4 | ipv6 } unicast

8. maximum-paths [ibgp] number-of-paths

9. (Optional) show running-config bgp

10. (Optional) copy running-config startup-config

upvoted 1 times

PimplePooper 11 months, 2 weeks ago

Selected Answer: B

B is correct

upvoted 1 times

jarz 1 year, 2 months ago

Selected Answer: C

Def. C. It's in this article from Cisco.
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/mpls/configuration/guide/mpls_cg/mp_vpn_multipath.html

upvoted 1 times

wts 1 year, 3 months ago

Selected Answer: C

First of all, the correct answer is: maximum-paths eibgp 2 . But it's not here.

You have to choose between A and B. Since the iBGP(M-BGP) is deeper in the MPLS, I choose to C.

upvoted 1 times

quyle 1 year, 2 months ago

P and PE using iBGP, PE and CE using eBGP. I think A B C is also correct. I don't know :)

upvoted 1 times

M_Abdulkarim 1 year, 4 months ago

Selected Answer: A

correct is A ==> eibgp

upvoted 1 times

cisconut 1 year, 5 months ago

Selected Answer: A

I could have chosen "B" until I read this "https://www.cisco.com/c/en/us/td/docs/ios/12_2sx/feature/guide/fsxeibmp.html#wp1027265". In Cisco document, it says "A" is the answer.

upvoted 1 times

Pbshah 1 year, 5 months ago

Selected Answer: B

This question does not tell which BGP load-balance it wants - whether it's iBGP or eBGP

upvoted 1 times

xziomal9 1 year, 7 months ago

Selected Answer: B

Sorry, B is correct

upvoted 1 times

xziomal9 1 year, 7 months ago

Selected Answer: D

Correct: D

upvoted 1 times

Question #92

Topic 1

Refer to the exhibit. After applying IPsec, the engineer observed that the DMVPN tunnel went down, and both spoke-to-spoke and hub were not establishing.

Which two actions resolve the issue? (Choose two.)

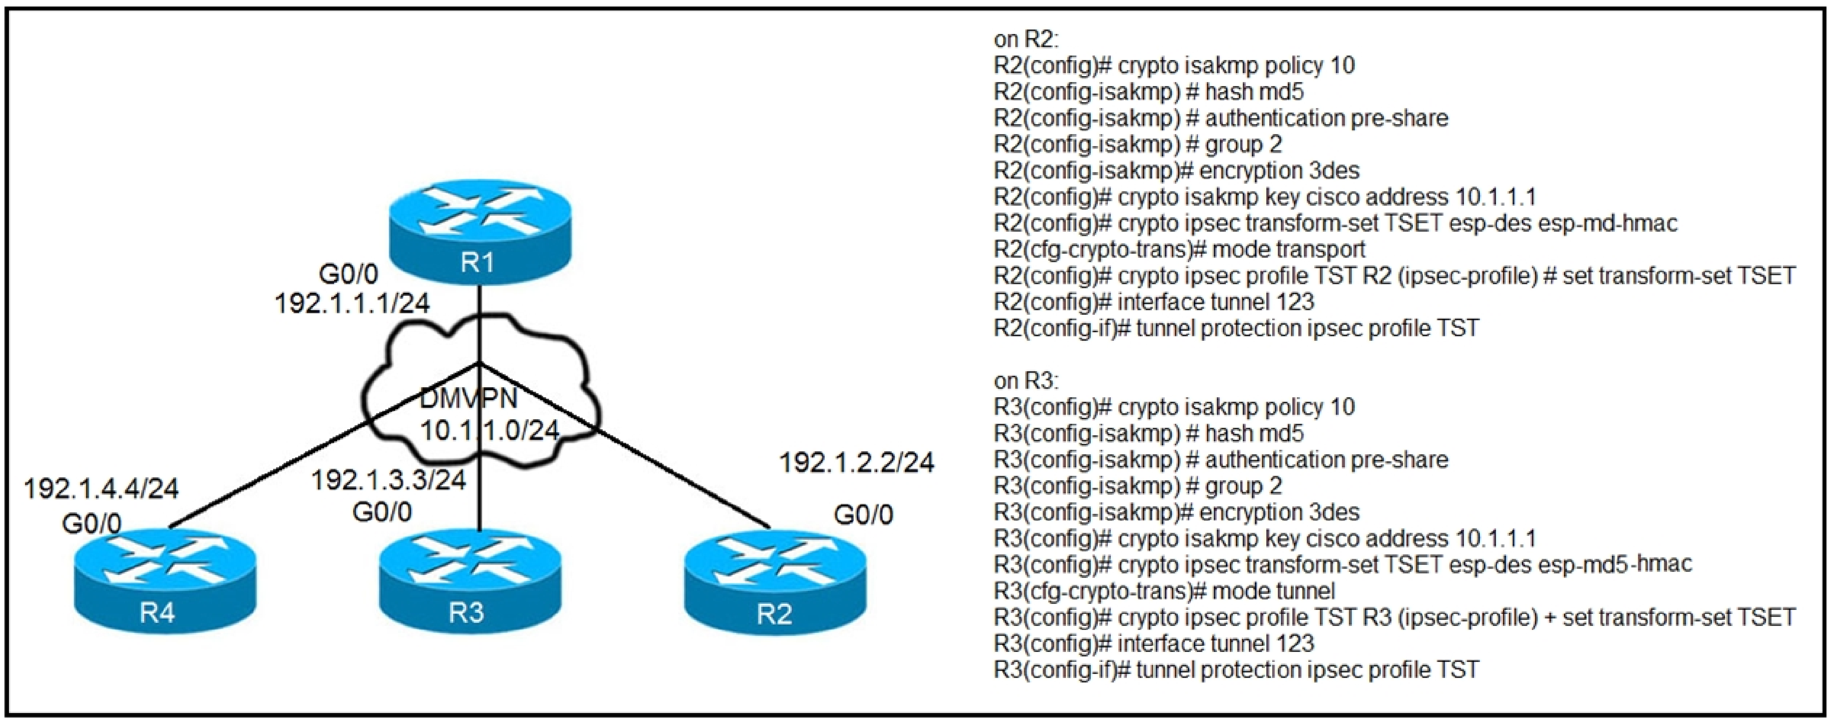

```
on R2:
R2(config)# crypto isakmp policy 10
R2(config-isakmp) # hash md5
R2(config-isakmp) # authentication pre-share
R2(config-isakmp) # group 2
R2(config-isakmp)# encryption 3des
R2(config)# crypto isakmp key cisco address 10.1.1.1
R2(config)# crypto ipsec transform-set TSET esp-des esp-md-hmac
R2(cfg-crypto-trans)# mode transport
R2(config)# crypto ipsec profile TST R2 (ipsec-profile) # set transform-set TSET
R2(config)# interface tunnel 123
R2(config-if)# tunnel protection ipsec profile TST

on R3:
R3(config)# crypto isakmp policy 10
R3(config-isakmp) # hash md5
R3(config-isakmp) # authentication pre-share
R3(config-isakmp) # group 2
R3(config-isakmp)# encryption 3des
R3(config)# crypto isakmp key cisco address 10.1.1.1
R3(config)# crypto ipsec transform-set TSET esp-des esp-md5-hmac
R3(cfg-crypto-trans)# mode tunnel
R3(config)# crypto ipsec profile TST R3 (ipsec-profile) + set transform-set TSET
R3(config)# interface tunnel 123
R3(config-if)# tunnel protection ipsec profile TST
```

A. Change the mode from mode tunnel to mode transport on R3.

B. Remove the crypto isakmp key cisco address 10.1.1.1 on R2 and R3.

C. Configure the crypto isakmp key cisco address 192.1.1.1 on R2 and R3.

D. Configure the crypto isakmp key cisco address 0.0.0.0 on R2 and R3.

E. Change the mode from mode transport to mode tunnel on R2.

Correct Answer: *AD*

Community vote distribution

AD (56%) BD (44%)

Guitarman Highly Voted 3 years, 3 months ago

I LITERALLY just labbed this. Please forgive the long explanation but I want to share for future testers. I was torn between changing the tunnel mode or removing one address and adding the other. B and D are definitely correct. You can't just put in the command with 0.0.0.0. If you do, you will end up with two crypto key commands and both addresses so the one to the tunnel address MUST be removed. Again, NO DOUBT...B AND D!!!!!

upvoted 21 times

spiderconnard 2 years, 4 months ago

Having many crypto keys is not an issue. you can leave the 10.1.1.1. If you add on top of it either 0.0.0.0 or 192.1.1.1 the tunnel protocol will go up.

upvoted 9 times

vdsdrs 2 years, 4 months ago

Does it mean that C and D are correct?

upvoted 2 times

T_Cos Most Recent 2 days, 5 hours ago

Options A and D are correct

upvoted 1 times

louisvuitton12 1 month, 3 weeks ago

Selected Answer: AD

Worked at Cisco TAC VPN team for over a year. A and D are correct.

upvoted 3 times

Ll123123 1 month, 4 weeks ago

Selected Answer: AD

AD I would say

upvoted 1 times

mouin 3 months, 1 week ago

Selected Answer: AD

I've been playing around with this lab for an hour.
The correct answer with no doubt is AD

upvoted 1 times

Brand 3 months, 2 weeks ago

Selected Answer: AD

I tested this scenario in my DMVPN lab just now. For this lab I configured ipsec transform-set with "mode tunnel" on hub and also in spokes. DMVPN was up, EIGRP was working etc. But than I changed the transform-set in spoke2 to "mode transport" and shut/no shut the tunnel interface. Spoke2 is not able to ping hub or the other spoke after that. So I'm 100% sure that one of the answers is "A" and looks like having multiple keys is not an issue so I'd go with "D" as well.

But before taking my comment as absolutely correct, lab it yourself.

upvoted 1 times

inteldarvid 4 months, 2 weeks ago

Selected Answer: BD

100 % B and D I check in lab

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: BD

Corerct Band D: You can't just put in the command with 0.0.0.0. If you do, you will end up with two crypto key commands and both addresses so the one to the tunnel address MUST be removed.

upvoted 1 times

Malasxd 7 months, 1 week ago

I LAB it and "D" is definilly right.

I didn't need to remove the old command. You can have many isakmp keys and it worked with the ends in different modes (I got surprised with that). So, I am not sure about the another right answer.

upvoted 2 times

Wooker 9 months ago

Selected Answer: BD

B and D correct

upvoted 1 times

Noproblem22 1 year, 1 month ago

BD makes more sense

upvoted 1 times

wts 1 year, 3 months ago

Seems like it can't be removed.

```
R2#show crypto isakmp key
Keyring Hostname/Address Preshared Key
default 8.8.8.8 cisco

Tunnel123 is up, line protocol is down

R2(config)#crypto isakmp key cisco address 0.0.0.0

Tunnel123 is up, line protocol is up

R2#show crypto isakmp key
Keyring Hostname/Address Preshared Key
default 8.8.8.8 cisco
0.0.0.0 [0.0.0.0] cisco
```

upvoted 2 times

xziomal9 1 year, 7 months ago

Selected Answer: BD

B and D

upvoted 2 times

Iarn 1 year, 7 months ago

Selected Answer: BD

10.1.1.1 is the inside Tunnel address the key need to be bond to the crypto key to either 192.1.1.1 or 0.0.0.0, Transport Modes dont matter

upvoted 2 times

Hack4 1 year, 10 months ago

A and D are the best

upvoted 1 times

Carl1999 1 year, 10 months ago

Selected Answer: AD

https://www.cisco.com/c/en/us/support/docs/security-vpn/dynamic-multi-point-vpn-dmvpn/116957-technote-dmvpn-00.html

upvoted 1 times

Carl1999 1 year, 10 months ago

B & D sorry
Communication is OK in both mode transport and tunnel, just for overhead issues.

upvoted 2 times

[Removed] 1 year, 10 months ago

Selected Answer: AD

With GRE the mode is already tunnel so need to specify, change R3 to tunnel mode and add D.

upvoted 2 times

[Removed] 1 year, 10 months ago

meant transport

upvoted 2 times

[Removed] 1 year, 10 months ago

R3 to transport

upvoted 1 times

[Removed] 1 year, 10 months ago

When using DMVPN with IPSec, it is unneccessary to use tunnel mode. Why? DMVPN uses GRE which means that a new IP header is already added by GRE. The GRE encapsulation happens on the tunnel interface before the encryption process takes place.

upvoted 1 times

Question #93 *Topic 1*

Which statement about route distinguishers in an MPLS network is true?

A. Route distinguishers allow multiple instances of a routing table to coexist within the edge router.

B. Route distinguishers are used for label bindings.

C. Route distinguishers make a unique VPNv4 address across the MPLS network.

D. Route distinguishers define which prefixes are imported and exported on the edge router.

Correct Answer: *C*

Community vote distribution

C (100%)

Ll123123 1 month, 4 weeks ago

C I would choose

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: C

correct anwser is C

upvoted 1 times

Hack4 1 year, 10 months ago

A is FALSE ..Route distinguishers allow multiple instances of a routing table to coexist within the edge router #### But allow multiple instances of the same IP PREFIX to coexist within the same router..

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct - Route distinguishers make a unique VPNv4 address across the MPLS network.

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

Sorry, the answer is A

upvoted 1 times

AliMo123 2 years, 4 months ago

A is the correct answer for VRF in general not for RD
C is the correct answer

upvoted 2 times

Hack4 1 year, 10 months ago

Nope the answer cannot be A ....A is FALSE .. It would be rather allow multiple instances of the same IP-PREFIX to coexist within the same router..

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

ZachTL11 2 years, 8 months ago

Answer is C.
VPNv4 prefix —IPv4 prefix preceded by an 8-byte route distinguisher. The VPN addresses are made unique by adding a route distinguisher to the front of the address.
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/15-mt/mp-l3-vpns-15-mt-book/mp-vpn-rte-target-rw.html

upvoted 1 times

Question #94 *Topic 1*

Which statement about MPLS LDP router ID is true?

A. If not configured, the operational physical interface is chosen as the router ID even if a loopback is configured.

B. The loopback with the highest IP address is selected as the router ID.

C. The MPLS LDP router ID must match the IGP router ID.

D. The force keyword changes the router ID to the specified address without causing any impact.

Correct Answer: *B*

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/12-4m/mp-ldp-12-4m-book.pdf

Community vote distribution

B (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

the anwser corret is B

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct - The loopback with the highest IP address is selected as the router ID.

upvoted 2 times

examShark 2 years, 4 months ago

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/12-4m/mp-ldp-12-4m-book.pdf

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

CraigB83 3 years, 2 months ago

B

"LDP will use the router ID as the source of TCP session. These IP has to be /32 and be in reachable. To modify the router ID use the command mpls ldp router-id (inteface) [force]. By default router ID is selected with the highest IP of a loopback, if non exist the highest IP of any active interfaces."

upvoted 4 times

Jack1188 3 years, 4 months ago

b win this question.

upvoted 1 times

Question #95

Topic 1

Refer to the exhibit. Which interface configuration must be configured on the spoke A router to enable a dynamic DMVPN tunnel with the spoke B router?

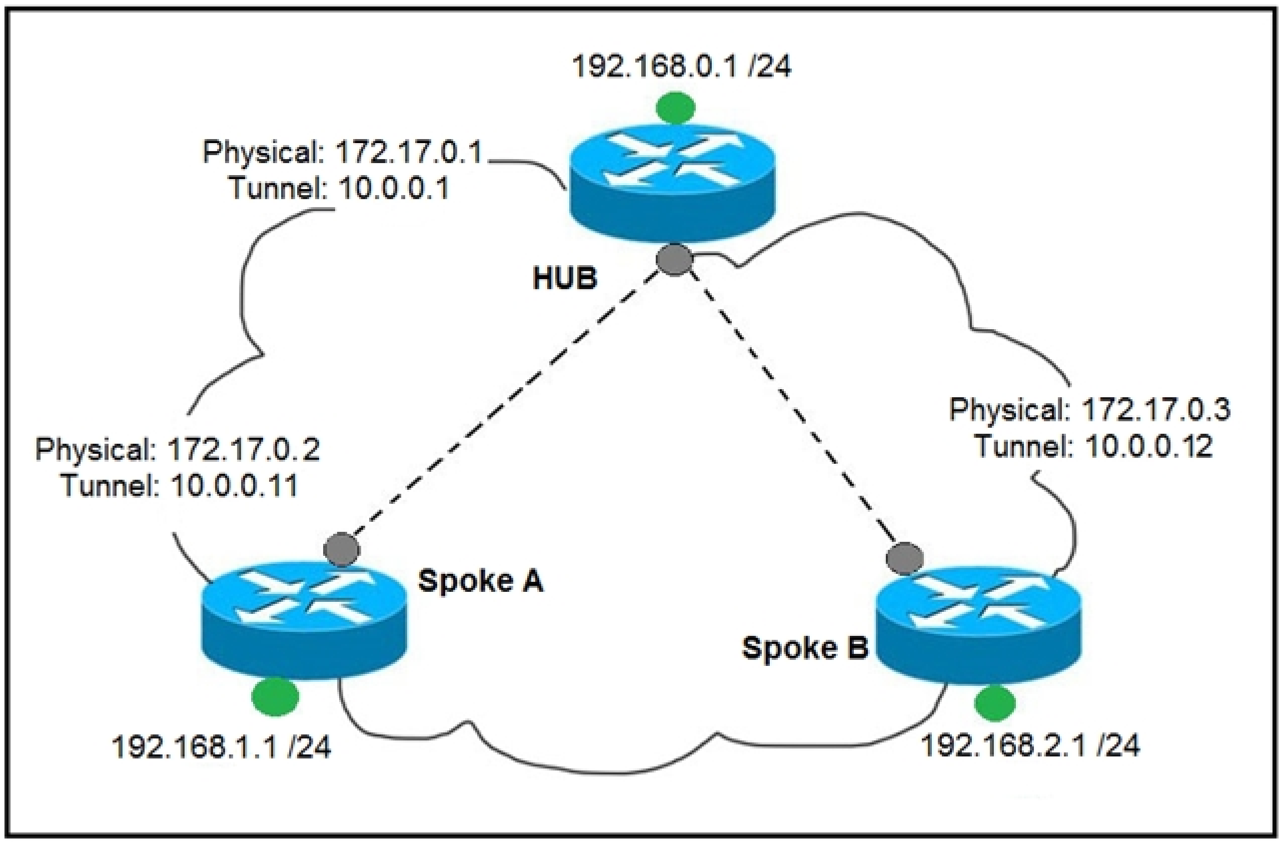

A.

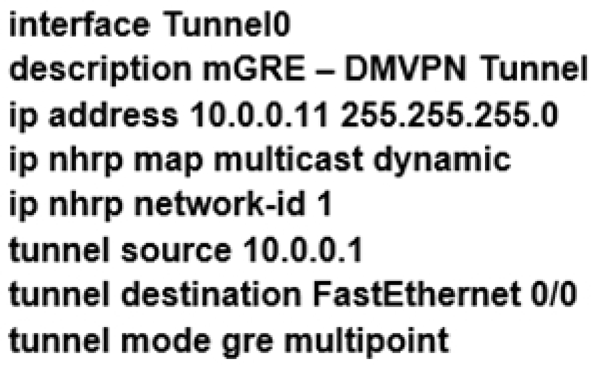

```
interface Tunnel0
description mGRE – DMVPN Tunnel
ip address 10.0.0.11 255.255.255.0
ip nhrp map multicast dynamic
ip nhrp network-id 1
tunnel source 10.0.0.1
tunnel destination FastEthernet 0/0
tunnel mode gre multipoint
```

B.

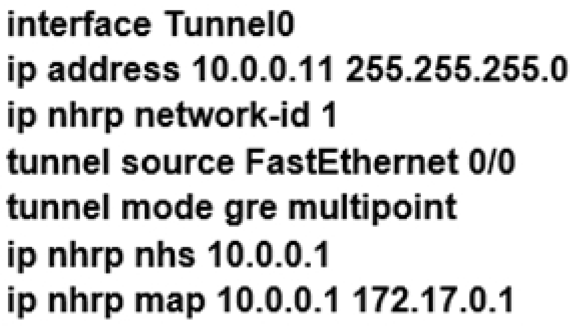

```
interface Tunnel0
ip address 10.0.0.11 255.255.255.0
ip nhrp network-id 1
tunnel source FastEthernet 0/0
tunnel mode gre multipoint
ip nhrp nhs 10.0.0.1
ip nhrp map 10.0.0.1 172.17.0.1
```

C.

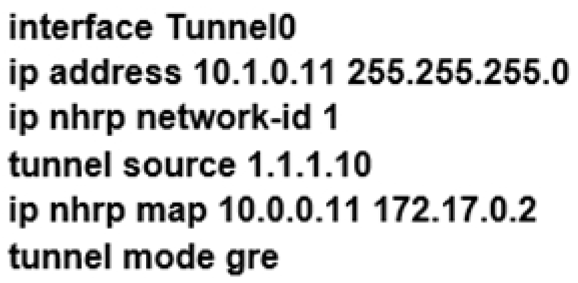

```
interface Tunnel0
ip address 10.1.0.11 255.255.255.0
ip nhrp network-id 1
tunnel source 1.1.1.10
ip nhrp map 10.0.0.11 172.17.0.2
tunnel mode gre
```

D.

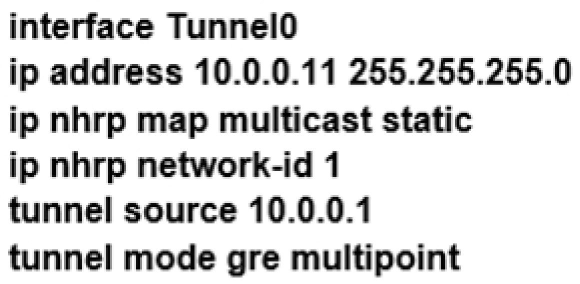

```
interface Tunnel0
ip address 10.0.0.11 255.255.255.0
ip nhrp map multicast static
ip nhrp network-id 1
tunnel source 10.0.0.1
tunnel mode gre multipoint
```

Correct Answer: *B*

⊟ **studybuddy10** Highly Voted 2 years, 1 month ago

B - only possible answer, all other sources are incorrect and B has correct NHRP map for the hub and NHS.

upvoted 5 times

⊟ **Colmenarez** Most Recent 3 months, 3 weeks ago

A is not correct, tunnel source can't be 10.0.0.1
B seems to be ok
C is not correct, missing multipoint on tunnel mode gre.
D is not correct, ip nhrp map multicast static is not a valid command.

The correct answer is B
upvoted 1 times

inteldarvid 5 months, 2 weeks ago
option B is correct
upvoted 1 times

HungarianDish 6 months, 2 weeks ago
Choosing "B" after all. - As any other answers are completely wrong. #ip nhrp map multicast x.x.x.x (NBMA of HUB) is still missing from "B".
upvoted 1 times

HungarianDish 8 months ago
It is hard to choose from these answers. Phase 2 would really need this on the spokes: #ip nhrp map multicast x.x.x.x
(#ip nhrp map multicast dynamic -> it is configured on the HUB, so obviously incorrect)
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_dmvpn/configuration/xe-16-7/sec-conn-dmvpn-xe-16-7-book/sec-conn-dmvpn-dmvpn.html
upvoted 1 times

Huntkey 1 year, 2 months ago
B will not allow spoke to spoke tunnel. ip nhrp map multicast dynamic and tunnel mode gre multipoint are required. I think the question would assume phase 1 already configured and working. Then would go with A to change it to phase 2
upvoted 1 times

Hack4 1 year, 10 months ago
B is correct .
upvoted 1 times

error_909 2 years, 3 months ago
The given answer is correct
upvoted 2 times

examShark 2 years, 4 months ago
The given answer is correct
Look at the tunnel source and destination
upvoted 2 times

Question #96 Topic 1

Which list defines the contents of an MPLS label?

A. 20-bit label; 3-bit traffic class; 1-bit bottom stack; 8-bit TTL

B. 32-bit label; 3-bit traffic class; 1-bit bottom stack; 8-bit TTL

C. 20-bit label; 3-bit flow label; 1-bit bottom stack; 8-bit hop limit

D. 32-bit label; 3-bit flow label; 1-bit bottom stack; 8-bit hop limit

Correct Answer: *A*

Reference:

https://tools.ietf.org/html/rfc5462

Community vote distribution

A (100%)

larn 1 year, 7 months ago

Selected Answer: A

32 Bit total, with 3 Bit COS

20-bit label; 3-bit traffic class; 1-bit bottom stack; 8-bit TTL = 32

upvoted 1 times

Hack4 1 year, 10 months ago

The given answer is correct

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct - 20-bit label; 3-bit traffic class; 1-bit bottom stack; 8-bit TTL

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

ZachTL11 2 years, 8 months ago

A is the correct answer.
In total the MPLS Label is 32 bits. (20+3+1+8)
Label: Label Value, 20 bits
Exp: Experimental Use, 3 bits
S: Bottom of Stack, 1 bit
TTL: Time to Live, 8 bits

upvoted 4 times

Question #97

Topic 1

Refer to the exhibit. What does the imp-null tag represent in the MPLS VPN cloud?

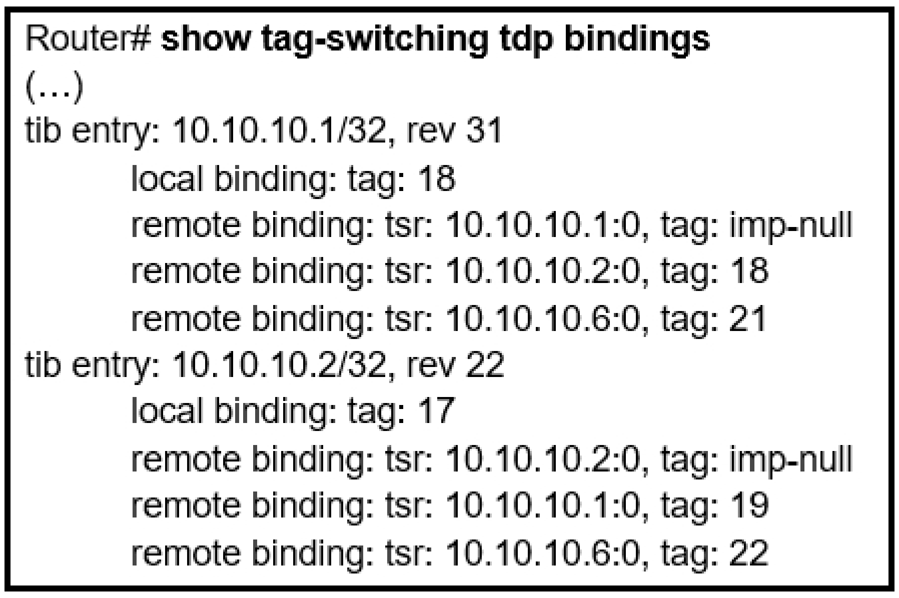

```
Router# show tag-switching tdp bindings
(…)
tib entry: 10.10.10.1/32, rev 31
        local binding: tag: 18
        remote binding: tsr: 10.10.10.1:0, tag: imp-null
        remote binding: tsr: 10.10.10.2:0, tag: 18
        remote binding: tsr: 10.10.10.6:0, tag: 21
tib entry: 10.10.10.2/32, rev 22
        local binding: tag: 17
        remote binding: tsr: 10.10.10.2:0, tag: imp-null
        remote binding: tsr: 10.10.10.1:0, tag: 19
        remote binding: tsr: 10.10.10.6:0, tag: 22
```

A. Pop the label

B. Impose the label

C. Include the EXP bit

D. Exclude the EXP bit

Correct Answer: *A*

Community vote distribution

A (100%)

anonymous1966 Highly Voted 3 years, 5 months ago

Just for reference:
Label-Switched Path (LSP) - The path that a labeled packet traverses through a network, from label imposition to disposition.

P/PE and C/CE -

P and PE routers are LSRs and LERs in the context of MPLS-VPN.

The term P comes from routers being in the provider network.
C routers are routers found in the customer network.

CE routers are the routers on the customer edge facing the provider.
PE routers are provider edge routers, which connect to the CE routers.
CE routers normally run plain IP (not required to be MPLS-aware).

Ref: https://www.ccexpert.us/traffic-engineering/mpls-terminology.html

upvoted 7 times

Dead_Adriano 2 years ago

This looks more like a comment to Q57 :-)

upvoted 1 times

adeeb1988ly Most Recent 9 months, 1 week ago

Answer A
Explanation:
The imp-null (implicit null) tag instructs the upstream router to pop the tag entry off the tag stack before forwarding the packet. Note: pop means remove the top MPLS labe

upvoted 2 times

Carl1999 1 year, 10 months ago

Selected Answer: A

PHP (Penultimate Hop Popping)

When exchanging labels, the LER informs the LSR of Implicit Null with a label value of 3.
When forwarding Implicit Null (imp-null) to the other party, LSR removes the label and forwards it.

upvoted 2 times

error_909 2 years, 3 months ago

The given answer is correct - Pop the label

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct.

TDP (the Tag Distribution Protocol)
https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/12492-mpls-tsh.html
https://www.ipspace.net/kb/tag/MPLS/Implicit_Explicit_NULL.html

upvoted 1 times

Question #98 *Topic 1*

DRAG DROP -

Drag and drop the MPLS terms from the left onto the correct definitions on the right.

Select and Place:

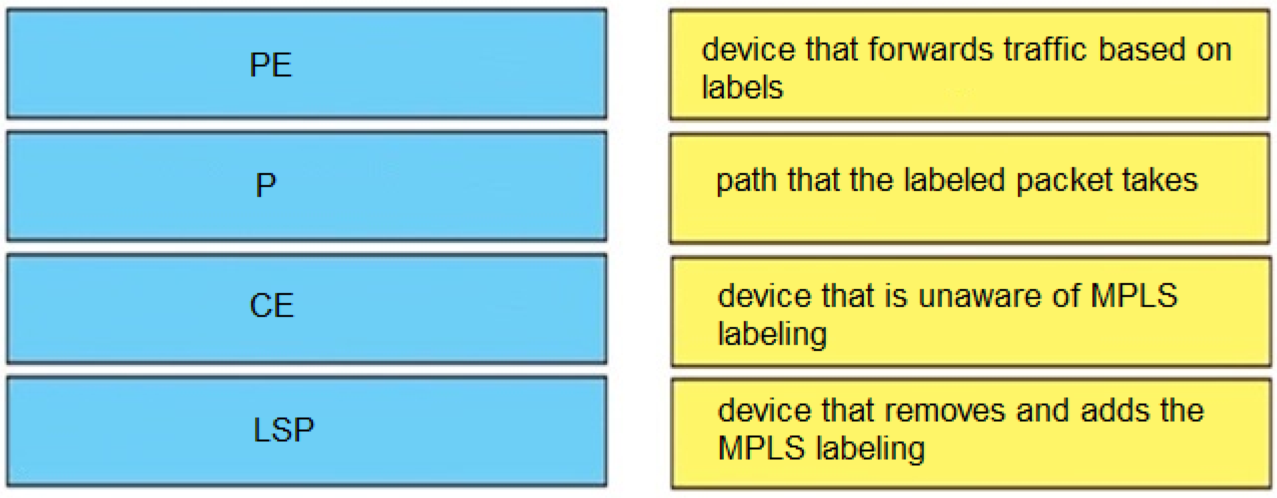

PE

P

CE

LSP

device that forwards traffic based on labels

path that the labeled packet takes

device that is unaware of MPLS labeling

device that removes and adds the MPLS labeling

Correct Answer:

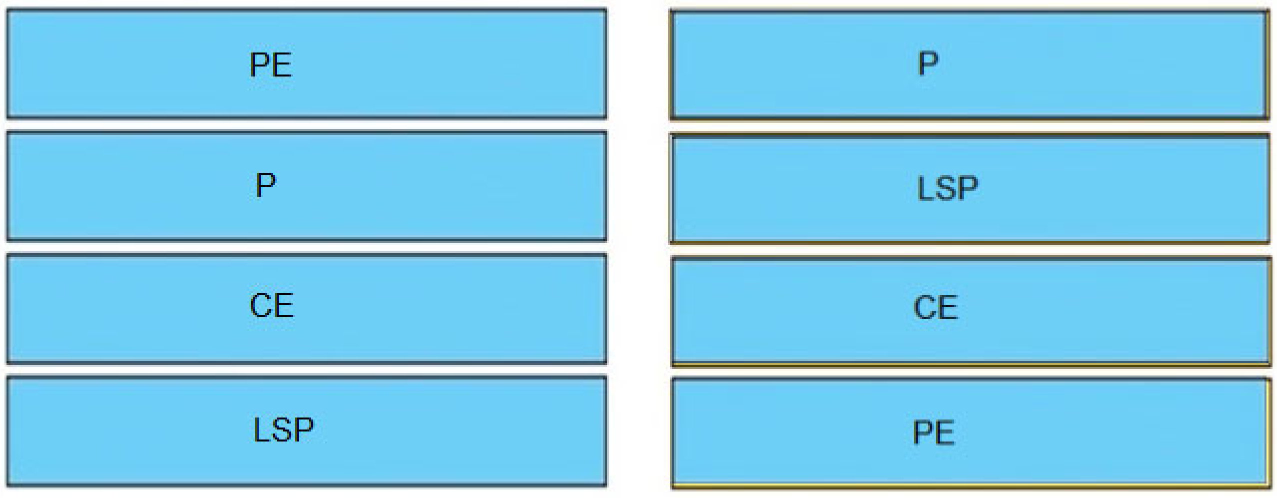

PE	P
P	LSP
CE	CE
LSP	PE

Dave513 Highly Voted 3 years ago

PE - Device that removes & adds the MPLS labelling;
P - Device that forwards traffic based on labels;
CE - Device that is unaware of MPLS labelling;
LSP - Path that labelled packets takes.

upvoted 5 times

error_909 Most Recent 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #99 Topic 1

Which transport layer protocol is used to form LDP sessions?

A. UDP

B. SCTP

C. TCP

D. RDP

Correct Answer: *C*

Community vote distribution

C (100%)

ZachTL11 Highly Voted 2 years, 8 months ago

Key word here is 'sessions'
UDP is connectionless and can be ruled out.
TCP is the answer.
The other two are just there to throw you off. RDP? really?

upvoted 5 times

wts Most Recent 1 year, 10 months ago

Selected Answer: C

"LDP uses User Datagram Protocol (UDP) and TCP to transport the protocol data unit (PDU) that carries LDP messages"

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Jack1188 3 years, 4 months ago

c win this question.

upvoted 1 times

Question #100 *Topic 1*

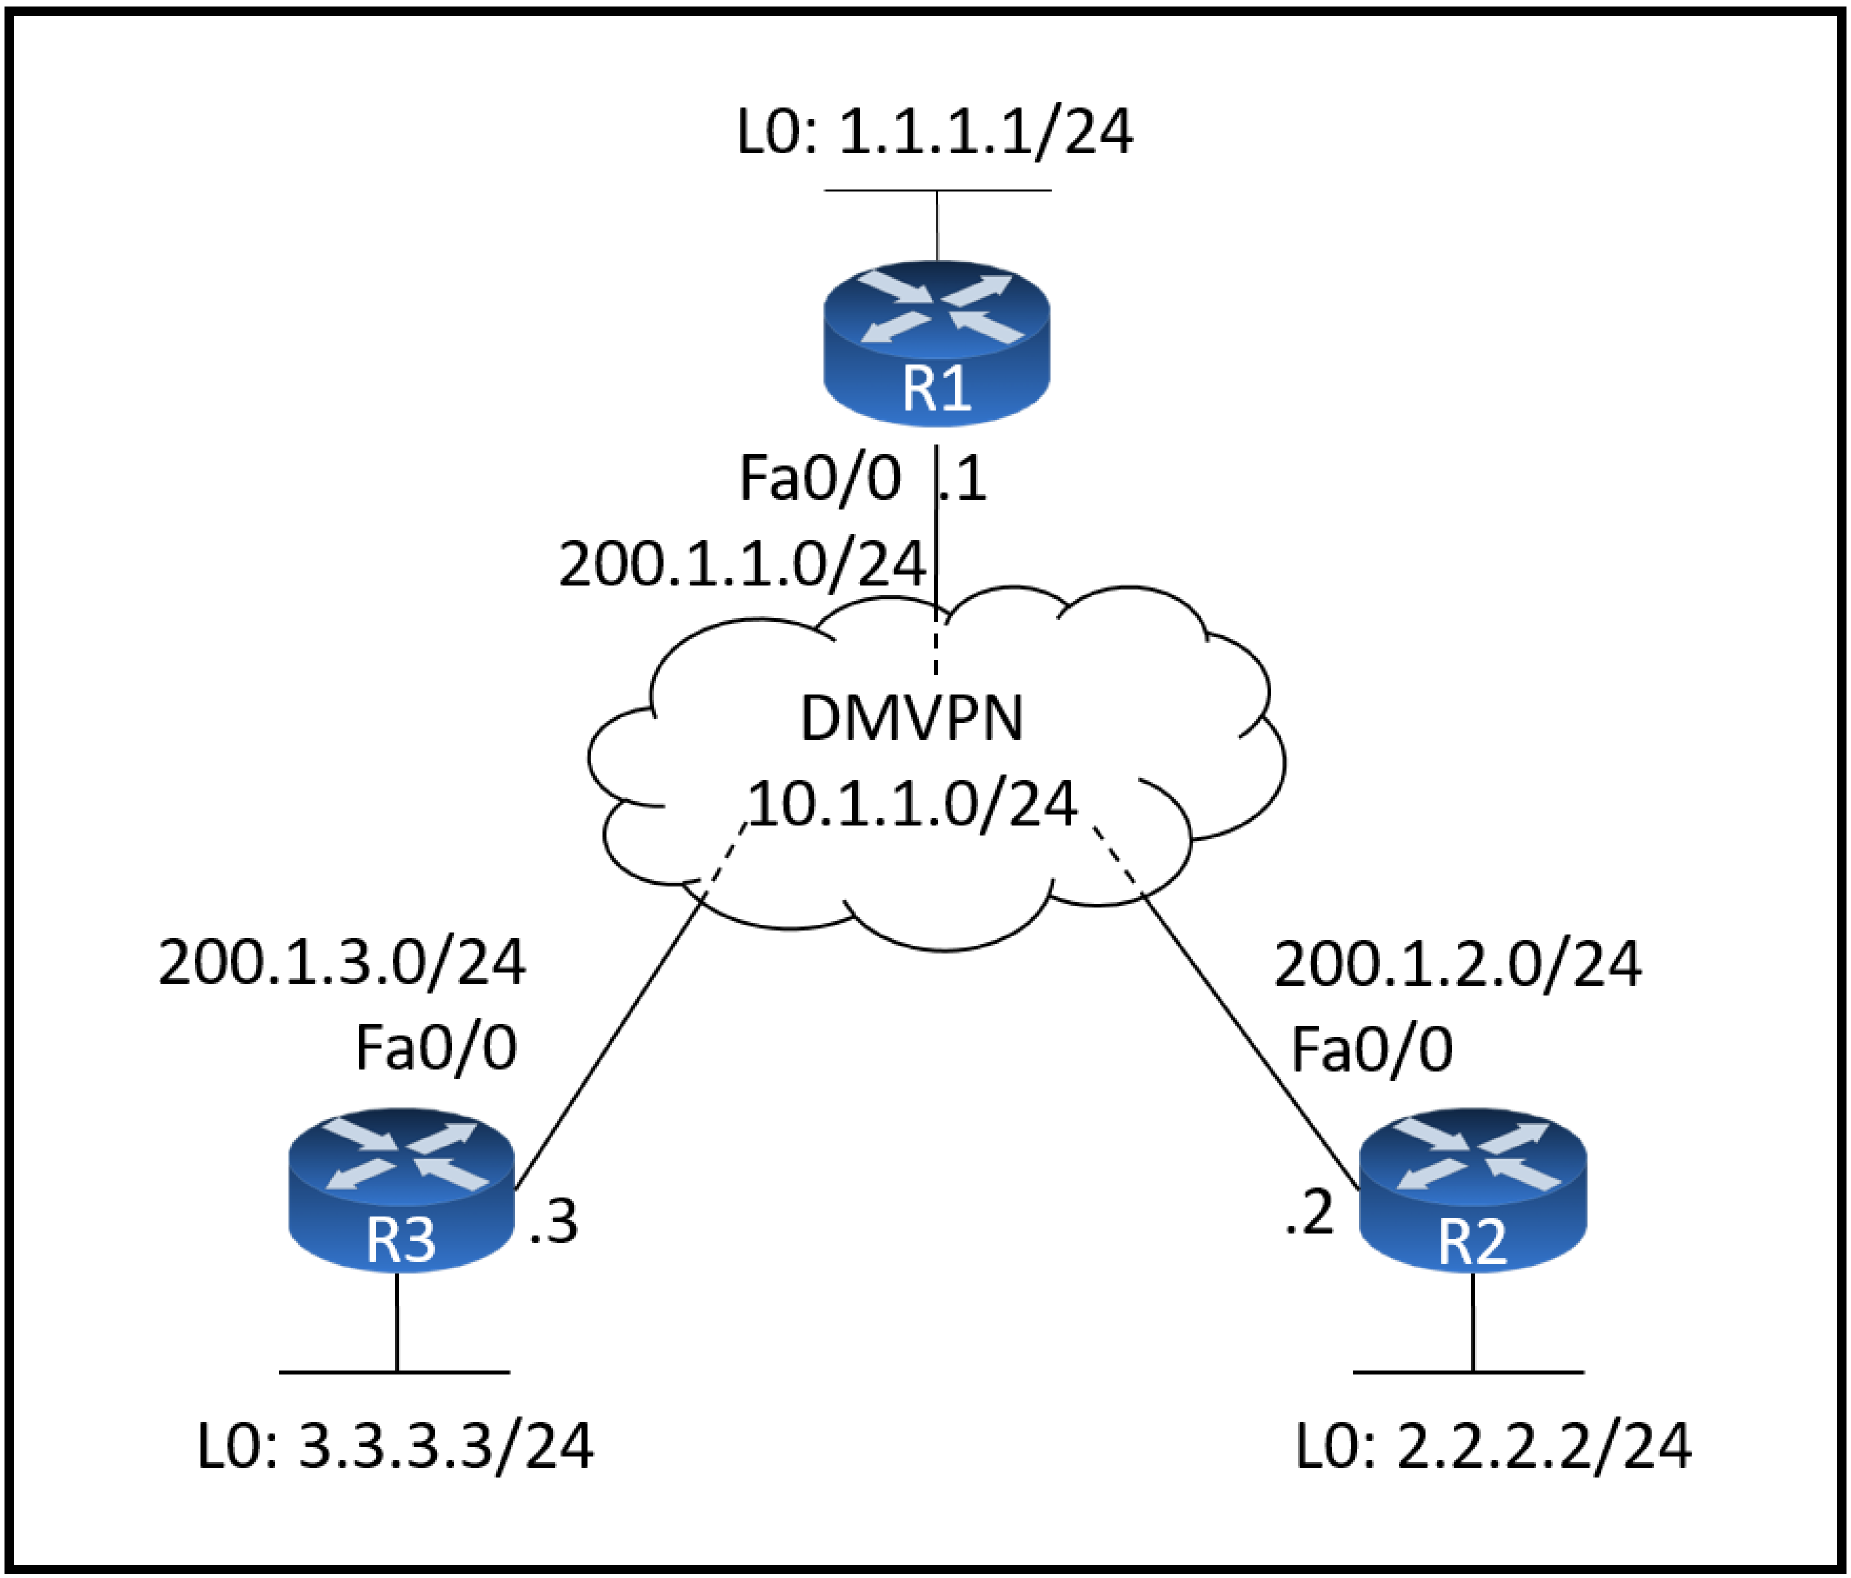

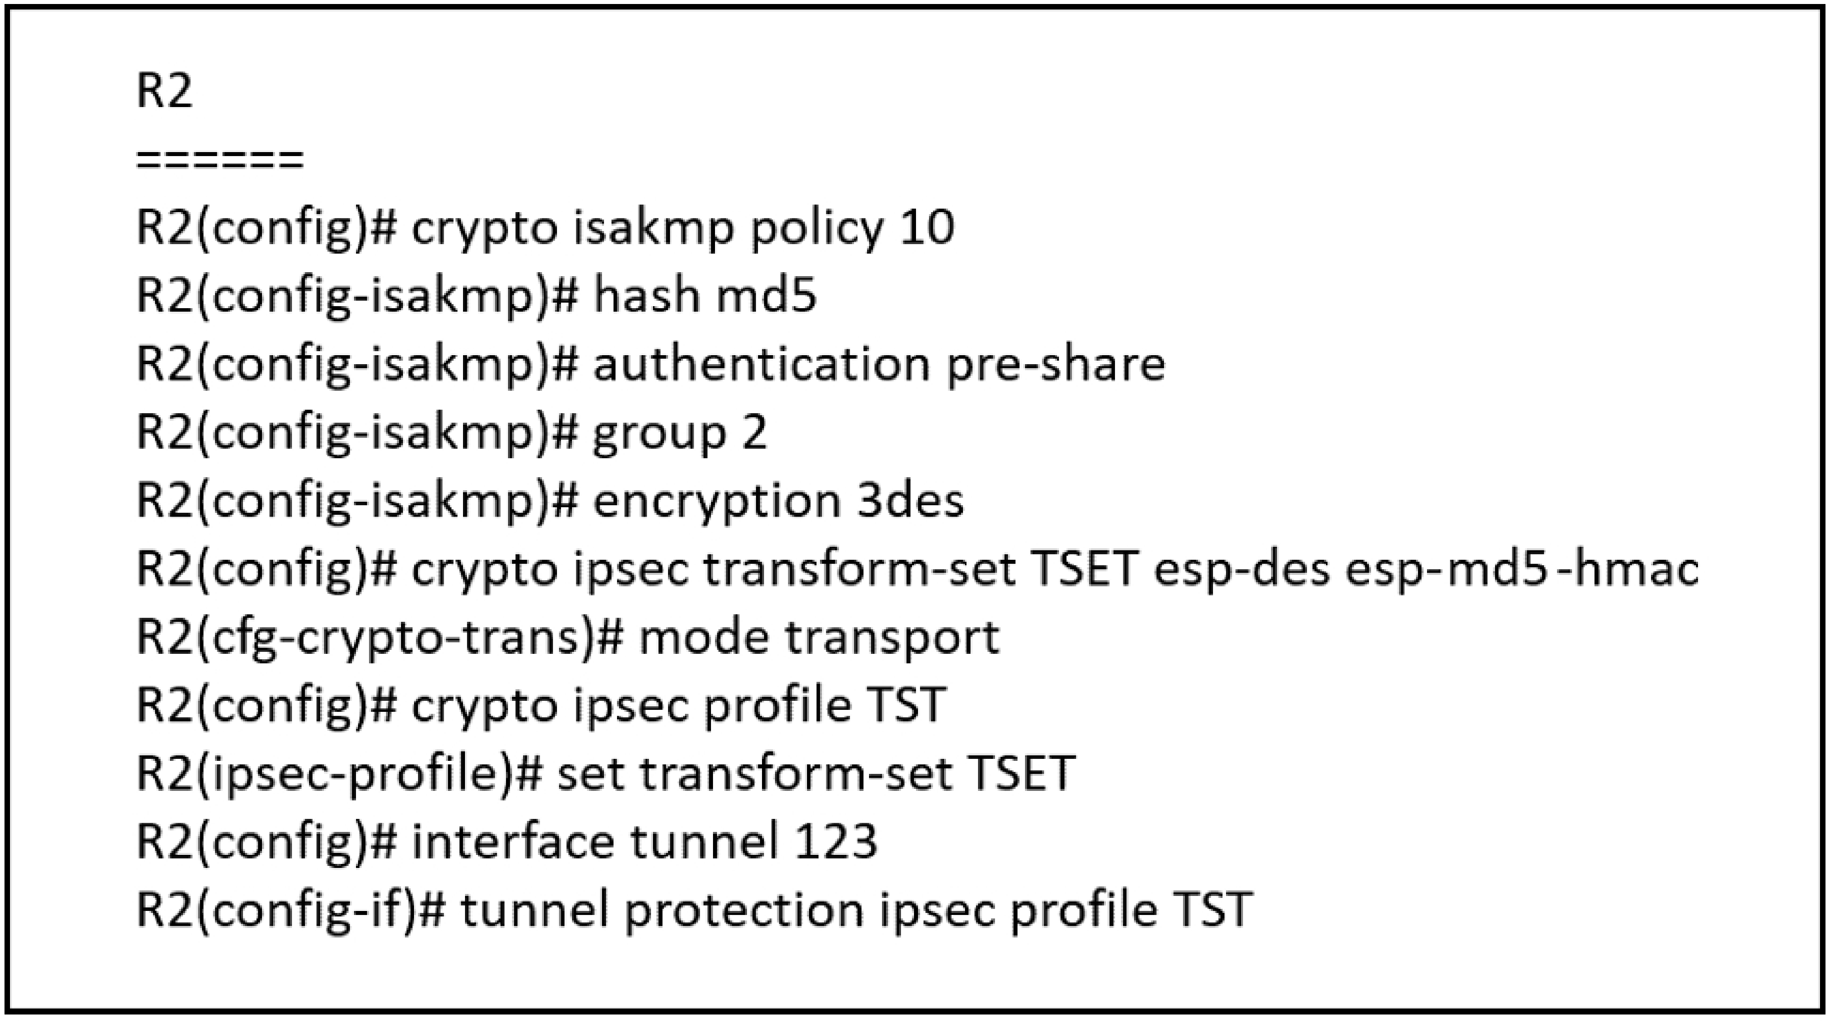

```
R2
======
R2(config)# crypto isakmp policy 10
R2(config-isakmp)# hash md5
R2(config-isakmp)# authentication pre-share
R2(config-isakmp)# group 2
R2(config-isakmp)# encryption 3des
R2(config)# crypto ipsec transform-set TSET esp-des esp-md5-hmac
R2(cfg-crypto-trans)# mode transport
R2(config)# crypto ipsec profile TST
R2(ipsec-profile)# set transform-set TSET
R2(config)# interface tunnel 123
R2(config-if)# tunnel protection ipsec profile TST
```

Refer to the exhibits.

Which configuration allows spoke-to-spoke communication using loopback as a tunnel source?

A. Configure crypto isakmp key cisco address 0.0.0.0 on the hub

B. Configure crypto isakmp key cisco address 200.1.0.0 255.255.0.0 on the hub

C. Configure crypto isakmp key cisco address 200.1.0.0 255.255.0.0 on the spokes

D. Configure crypto isakmp key cisco address 0.0.0.0 on the spokes

Correct Answer: *D*

Community vote distribution

D (100%)

HungarianDish 8 months ago

Selected Answer: D

for spoke-to-spoke we need to add this on the spokes too
https://community.cisco.com/t5/vpn/isakmp-with-0-0-0-0-dmvpn/td-p/4312380

upvoted 4 times

chris7890 1 year, 1 month ago

Is it possible that the command must be executed on the hub and on the spoke router?

Configure ISAKMP on all devices:

...

crypto isakmp key cisco address 0.0.0.0

https://ccieme.wordpress.com/2021/09/09/cisco-dynamic-multipoint-vpn/

upvoted 3 times

Bruffas 1 year, 9 months ago

I would assume that since we see the config on one spoke, that alternative A already is set on the HUB.
In that case D is the only answer that makes sense.

upvoted 3 times

studybuddy10 2 years, 1 month ago

Given answer is correct - D , the spokes dynamic tunnels with loopback sources are coming from 2.2.2.2 and 3.3.3.3 so only spokes with 0.0.0.0 would satisfy that.

upvoted 4 times

FrankZane 2 years, 1 month ago

I think A is correct
https://www.cisco.com/en/US/technologies/tk583/tk372/technologies_white_paper0900aecd802b8f3c.html

upvoted 1 times

Question #101 Topic 1

How does an MPLS Layer 3 VPN function?

A. multiple customer sites interconnect through service provider network to create secure tunnels between customer edge devices

B. multiple customer sites interconnect through a service provider network using customer edge to provider edge connectivity

C. set of sites interconnect privately over the Internet for security

D. set of sites use multiprotocol BGP at the customer site for aggregation

Correct Answer: *B*

Community vote distribution

B (67%) A (33%)

gndrx78 Highly Voted 1 year, 11 months ago

Given answer is ok. Reference: https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k-r6-5/lxvpn/configuration/guide/b-l3vpn-cg-asr9000-65x/b-l3vpn-cg-asr9000-65x_chapter_010.pdf

upvoted 10 times

MasoudGhorbani Most Recent 1 month, 2 weeks ago

B is the right answer because MPLS traffic is already going through secure ISP routers, not the internet. The main focus in MPLS is on segregating routes.

upvoted 1 times

JeffJeffson 5 months, 1 week ago

Selected Answer: B

Connectivity is between CE and PE in a provider environment.

upvoted 1 times

MicMillon 5 months, 3 weeks ago

Selected Answer: B

MPLS tunnels the routes through the providers core, but doesn't extend that tunnel to the edge device

upvoted 1 times

HungarianDish 6 months, 4 weeks ago

Selected Answer: B

Customer edge simply provides edge connectivity for the customer site. CE is not part of the provider mpls network (LSP).
https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2020/pdf/DGTL-BRKMPL-1100.pdf

upvoted 4 times

Malasxd 7 months, 2 weeks ago

I think it's "A". MPLS Layer 3 VPN is not secure. We don't configure any mechanism to make the tunnel secure like a ipsec for exemple.

And I've never seen documentation mentioning the word "security" for this type of tunnel.

upvoted 1 times

Malasxd 7 months, 1 week ago

I meant "B"

upvoted 2 times

6dd4aa0 8 months, 3 weeks ago

Selected Answer: A

Referring to Figure 18-11 in Page 740 of the CCNP Enterprise Advanced Routing ENARSI Official Guide, it shows that multiple CE routers are connected to a single PE as an ingress router. A VPN tunnel is formed from the ingress routers to a series of P (provider) routers to the egress PE routers before exiting the respective customer's site. Hence, the answer I will take is A

upvoted 3 times

Question #102 *Topic 1*

DRAG DROP -

Drag and drop the LDP features from the left onto the descriptions on the right.

Select and Place:

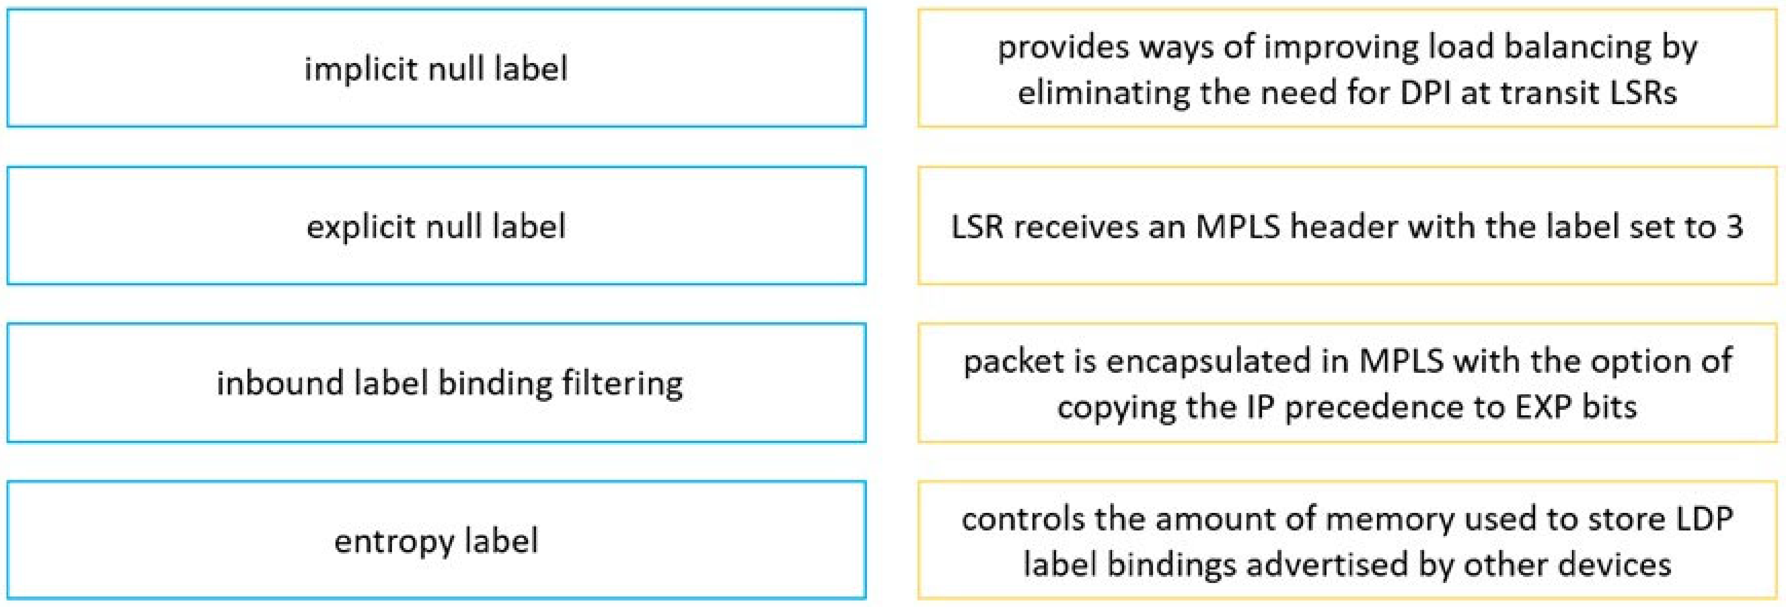

implicit null label

explicit null label

inbound label binding filtering

entropy label

provides ways of improving load balancing by eliminating the need for DPI at transit LSRs

LSR receives an MPLS header with the label set to 3

packet is encapsulated in MPLS with the option of copying the IP precedence to EXP bits

controls the amount of memory used to store LDP label bindings advertised by other devices

Correct Answer:

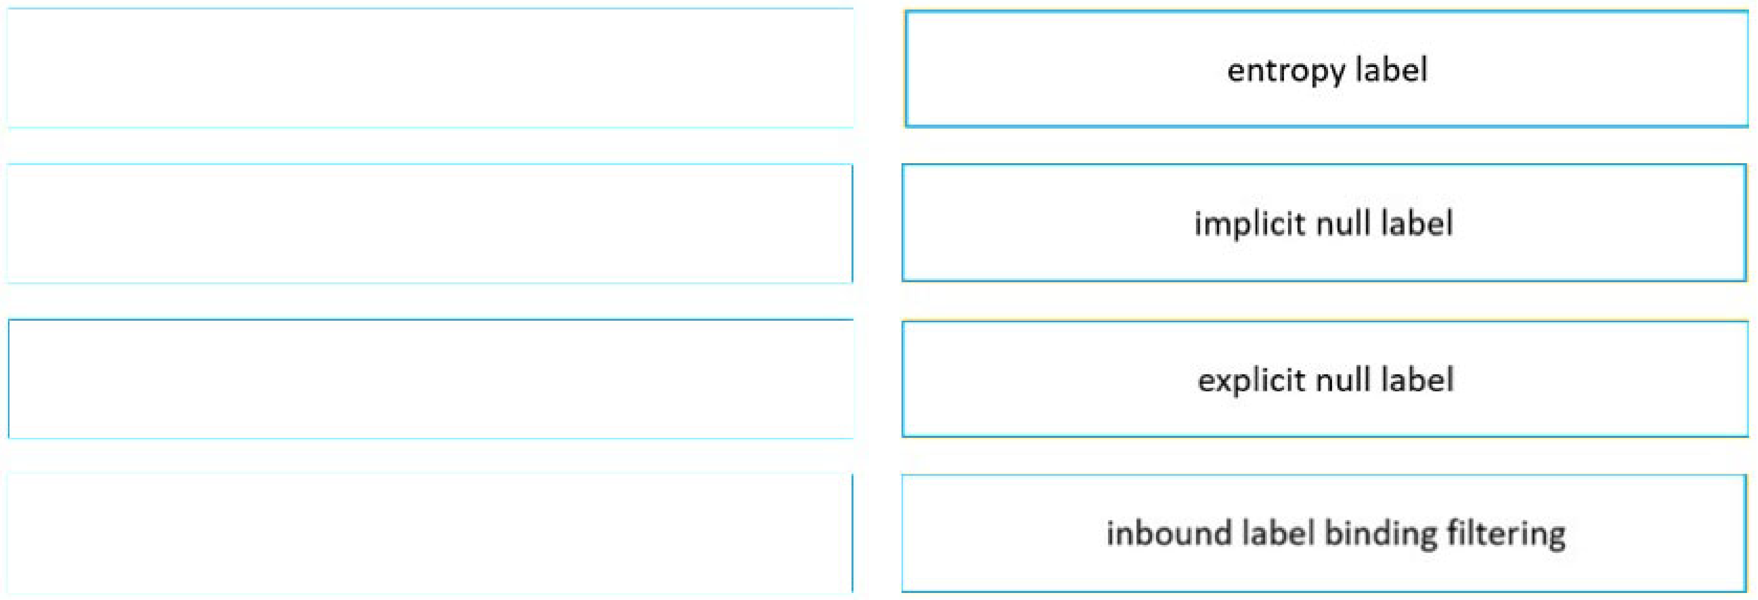

entropy label

implicit null label

explicit null label

inbound label binding filtering

HungarianDish 8 months ago

Correct.
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/xe-16-6/mp-ldp-xe-16-6-book/mp-ldp-inbound-filtr.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/xe-16/mp-ldp-xe-16-book/mp-ldp-entropy.html
https://www.ipspace.net/kb/tag/MPLS/Implicit_Explicit_NULL.html

upvoted 3 times

Question #103

Topic 1

Which two protocols work in the control plane of P routers across the MPLS cloud? (Choose two.)

A. ECMP

B. LDP

C. RSVP

D. MPLS OAM

E. LSP

Correct Answer: *BC*

Community vote distribution

BC (100%)

⊟ **HungarianDish** 8 months ago

Selected Answer: BC

Correct.

https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2020/pdf/DGTL-BRKMPL-1100.pdf

upvoted 3 times

Question #104 *Topic 1*

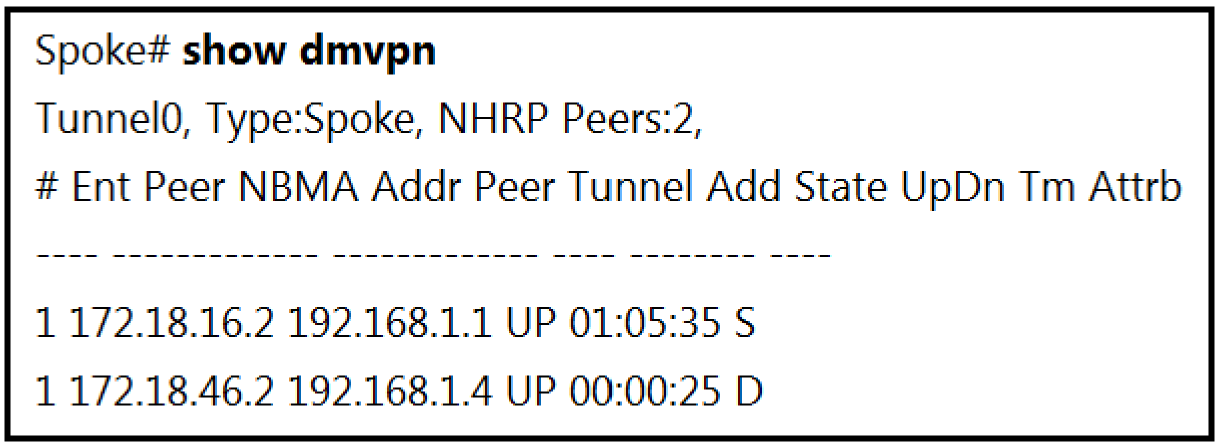

```
Spoke# show dmvpn
Tunnel0, Type:Spoke, NHRP Peers:2,
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb
---- ------------- ------------- ---- -------- ----
1 172.18.16.2 192.168.1.1 UP 01:05:35 S
1 172.18.46.2 192.168.1.4 UP 00:00:25 D
```

Refer to the exhibit. An engineer has configured DMVPN on a spoke router.
What is the WAN IP address of another spoke router within the DMVPN network?

A. 172.18.46.2

B. 172.18.16.2

C. 192.168.1.1

D. 192.168.1.4

Correct Answer: *D*

Community vote distribution

A (100%)

DaanB Highly Voted 2 years, 8 months ago

The WAN IP is 172.18.46.2

upvoted 15 times

Dave22 Highly Voted 2 years, 7 months ago

Answer is A please update

upvoted 5 times

Dave22 2 years, 7 months ago

This reason being is the NMBA address is the private and the Peer is the public and the first one which is 172.18.16.2 is not is because that is the Hub configured with a static mapping as it has "S" on the end. However when the address 172.18.46.2 is learned by the spoke it places a "D" at the end

upvoted 3 times

Calyfas Most Recent 10 months ago

A is correct

upvoted 1 times

Lilienen 10 months, 2 weeks ago

Selected Answer: A

A is correct

upvoted 1 times

Koume 11 months ago

Selected Answer: A

The option A , first the NMBA address are the wan ip address and second the NHRP registration is marked as dynamic (D). registration made to the hub are marked as static (S)

upvoted 1 times

Hurk2 11 months, 2 weeks ago

Selected Answer: A

A is correct, the D flag is dynamic so that is the spoke, the hub has the S (static) flag

upvoted 1 times

ChillingAgain 1 year, 1 month ago

Selected Answer: A

Answer is A
IP-Address 192.168.1.4 is the tunnel interface address
IP-Address 172.18.46.2 is the WAN interface address
D stands for dynamic tunnel which is created between spokes

upvoted 1 times

TECH3K3 1 year, 5 months ago

Selected Answer: A

Answer is A.
The WAN IP is NBNA address, which is the device physical interface.
S will be the HUB.
D will be he spoke who learn the other spoke address dynamically.
When you do the command "show dmvpn". The first IP address is the NBMA address and the IP address after is the Tunnel IP address.

upvoted 2 times

larn 1 year, 7 months ago

Selected Answer: A

Cannot be the S = Static As that will be the Hub & Peer NBMA Addr is the Outside/WAn address

upvoted 1 times

timtgh 1 year, 6 months ago

Other way around. Hubs are Dynamic, so S (Static) means it's a spoke.

upvoted 1 times

timtgh 1 year, 6 months ago

Disregard. Spoke-to-hub connection is S when seen on spoke side.

upvoted 1 times

Icy1 1 year, 10 months ago

hard to say what cisco means by "WAN IP". It points more to "inner" IP than to "outer/NBMA" IP. If we stick to fact that "WAN" is customer's wide area network, then they ask about inner IP - which is D. I don't like question where they juggle with words...

upvoted 1 times

timtgh 1 year, 6 months ago

WAN address is physical address of WAN interface, not the logical address of the tunnel interface. Not sure where you see the word "inner" used. What are they juggling? Physical address is the 172.18.X.X address on the left, which is A.

upvoted 1 times

YaPet 1 year, 10 months ago

Selected Answer: A

Agree with others, A is correct

upvoted 1 times

Hack4 1 year, 10 months ago

A is correct

upvoted 1 times

tyh391 1 year, 10 months ago

Selected Answer: A

As in discussion

upvoted 1 times

tyh391 1 year, 10 months ago

Selected Answer: A

Answer is A

upvoted 1 times

JingleJangus 1 year, 10 months ago

Selected Answer: A

Correct answer is A. Its in the NBMA field in the output of the sh cmd. The S on the far right means that neighbor is the HUB, dynamic or D is other spokes because they are dynamically discovered.

upvoted 2 times

Alex147 1 year, 11 months ago

Selected Answer: A

Corect is A.
Attribute D - dynamic
Attribute S - static

upvoted 2 times

Dead_Adriano 2 years ago

"D" for "dynamic" which means another spoke.
NBMA address for WAN so answer is A

upvoted 1 times

Question #105 *Topic 1*

What are two functions of LDP? (Choose two.)

A. It advertises labels per Forwarding Equivalence Class.

B. It uses Forwarding Equivalence Class.

C. It is defined in RFC 3038 and 3039.

D. It requires MPLS Traffic Engineering.

E. It must use Resource Reservation Protocol.

Correct Answer: *AB*

Community vote distribution

AB (100%)

⊟ **Hack4** 1 year, 10 months ago

yes i agree

upvoted 1 times

⊟ **Budh** 1 year, 10 months ago

Selected Answer: AB

Answer is correct

upvoted 2 times

⊟ **error_909** 2 years, 3 months ago

The given answer is correct

upvoted 1 times

⊟ **examShark** 2 years, 4 months ago

The given answer is correct

upvoted 2 times

⊟ **yuri_yuri** 2 years, 8 months ago

The correct answer is A & B
https://www.ccexpert.us/mpls-network/forwarding-equivalence-class.html

upvoted 2 times

Question #106 *Topic 1*

DRAG DROP -

Drag and drop the operations from the left onto the locations where the operations are performed on the right.

Select and Place:

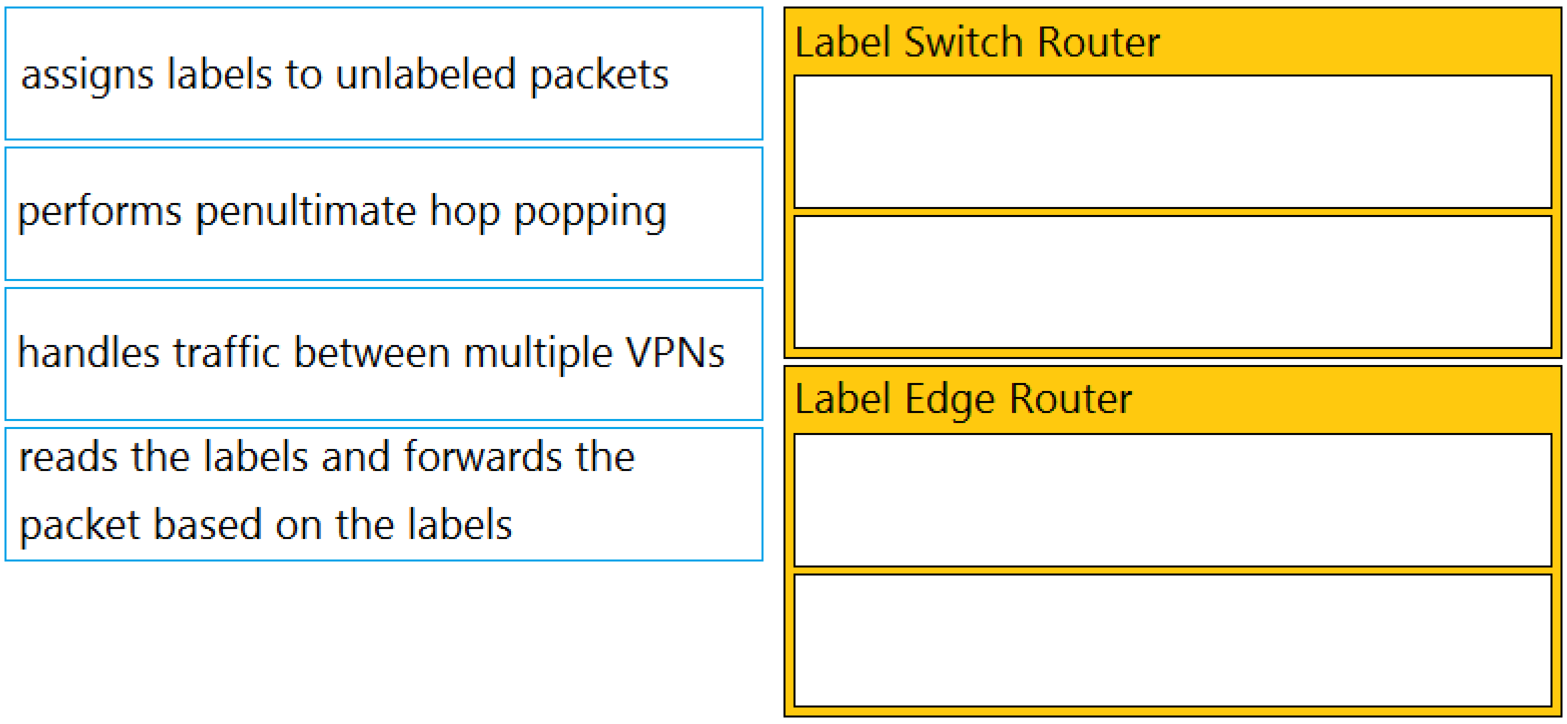

assigns labels to unlabeled packets

performs penultimate hop popping

handles traffic between multiple VPNs

reads the labels and forwards the packet based on the labels

Label Switch Router

Label Edge Router

Correct Answer:

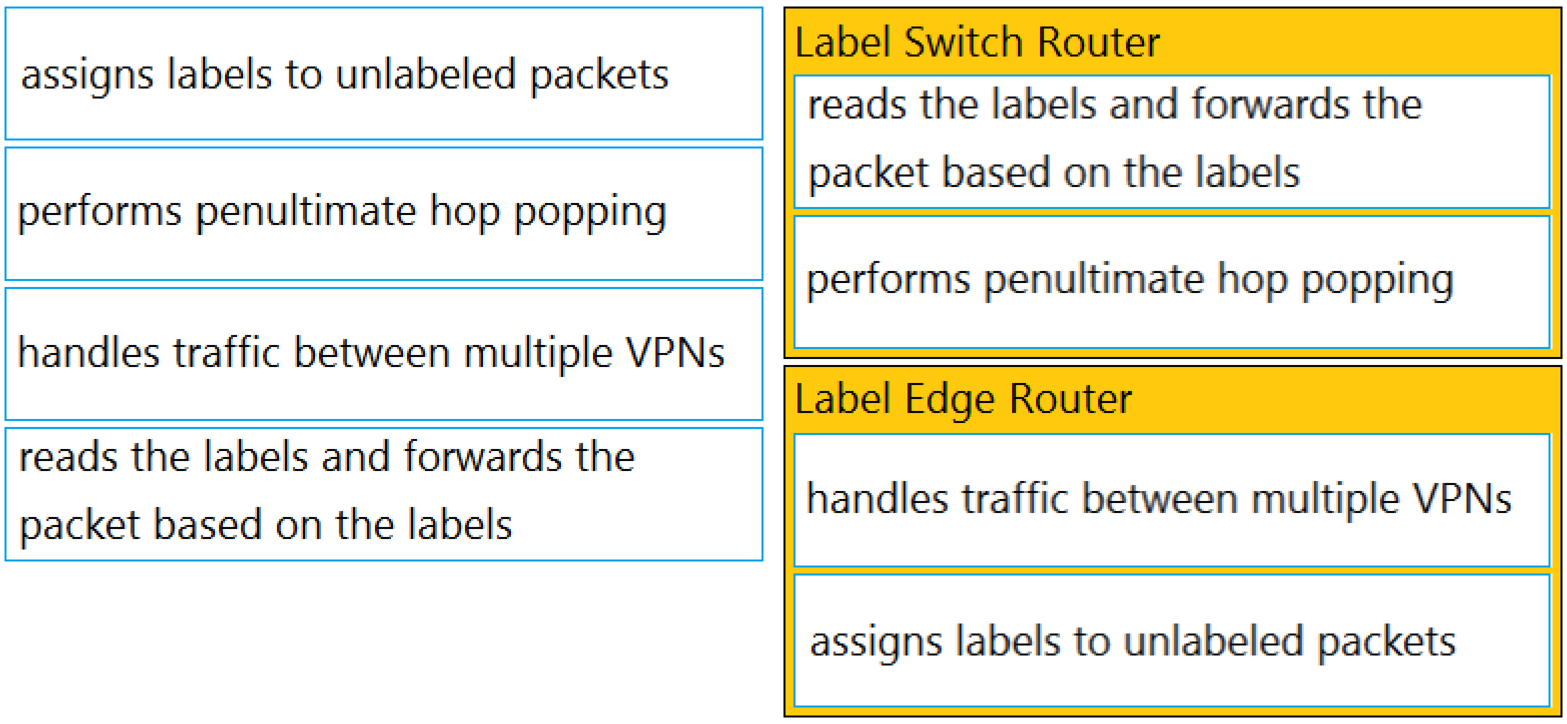

assigns labels to unlabeled packets

performs penultimate hop popping

handles traffic between multiple VPNs

reads the labels and forwards the packet based on the labels

Label Switch Router
- reads the labels and forwards the packet based on the labels
- performs penultimate hop popping

Label Edge Router
- handles traffic between multiple VPNs
- assigns labels to unlabeled packets

Chiaretta 5 months, 1 week ago

The given answer is correct. The PHP is made by LSR.

upvoted 2 times

Ash78 1 year, 8 months ago

Do they have to be in order? For example, can assigns labels handles change their places?

upvoted 2 times

timtgh 1 year, 6 months ago

Order doesn't matter within each box, just have to be in the right box.

upvoted 2 times

Hack4 1 year, 10 months ago

YES THE GIVEN ANSWER IS CORRECT

upvoted 2 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

Question #107 *Topic 1*

Which protocol does MPLS use to support traffic engineering?

A. TDP

B. RSVP

C. LDP

D. BGP

Correct Answer: *B*

GreatDane 1 year, 5 months ago

Ref: How MPLS Traffic Engineering works - Cisco Community

"...
These components work together to make MPLS TE work:
...
• Path setup is a signaling protocol to reserve the resources for a traffic flow and to establish the LSP for a traffic flow. Resource Reservation Protocol (RSVP) is used for this purpose and has been enhanced with TE extensions for carrying labels and building the LSP. An alternative to RSVP for MPLS TE is constrained routing with Label Distribution Protocol (LDP), but Cisco devices do not support this protocol.
..."

A. TDP

Wrong answer.

B. RSVP

Correct answer.

C. LDP

Wrong answer.

D. BGP

Wrong answer.

upvoted 2 times

Mjestic 2 years, 3 months ago

B is correct.
MPLS traffic engineering automatically establishes and maintains LSPs across the backbone by using RSVP.
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_path_setup/configuration/xe-16-11/mp-te-path-setup-xe-16-11-book/mpls-traffic-engineering-and-enhancements.html

upvoted 2 times

Question #108 *Topic 1*

An engineer configured a company's multiple area OSPF Head Office router and Site A Cisco routers with VRF lite. Each site router is connected to a PE router of an MPLS backbone:

```
Head Office & Site A -
ip cef
ip vrf abc
rd 101:101
!
interface FastEthernet0/0
ip vrf forwarding abc
ip address 172.16.16.X 255.255.255.252
!
router ospf 1 vrf abc
log-adjacency-changes
network 172.16.16.0 0.0.0.255 area 1
```

After finishing both site router configurations, none of the LSA 3, 4, 5, and 7 are installed at Site A router.
Which configuration resolves this issue?

A. configure capability vrf-lite on Site A and its connected PE router under router ospf 1 vrf abc

B. configure capability vrf-lite on both PE routers connected to Head Office and Site A routers under router ospf 1 vrf abc

C. configure capability vrf-lite on Head Office and its connected PE router under router ospf 1 abc

D. configure capability vrf-lite on Head Office and Site A routers under router ospf 1 vrf abc

Correct Answer: *D*

Community vote distribution

D (71%) A (29%)

myrmike Highly Voted 2 years ago

Notice that three of the answers involve configuring the PE router also. Since the engineer configured the company's router he presumably works for the company and not the ISP so the engineer would not have access to the PE router(s)

upvoted 16 times

guy276465281819372 Most Recent 4 months, 2 weeks ago

Selected Answer: D

D is right

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

https://community.cisco.com/t5/routing/where-to-configure-the-quot-capability-vrf-lite-quot-on-ce-or-pe/td-p/2812305

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

the answer corret is D:

https://forum.networklessons.com/t/when-and-where-to-use-capability-vrf-lite/14877

upvoted 1 times

Edwinmolinab 1 year, 4 months ago

Selected Answer: A

Answer: A

Explanation

In this case both Head Office and Site A routers run VRF (and OSPF) although they are CE routers.

So we must configure "capability vrf-lite" on them too.

For your information, the capability vrf-lite command disables the DN-bit (down bit) and domain-tag checks in OSPF. Since the CE router acts as the PE router in VRF-lite, these checks should be disabled, because the PE routers advertise VPN routes with DN-bit set to the CE routers. If the CE routers receive routes with DN-bit set, it will discard them. Hence, the checks should be disabled.

upvoted 4 times

GreatDane 1 year, 5 months ago

Selected Answer: D

Ref: Solved: Where to configure the "capability vrf lite", on CE or PE? - Cisco Community

Post by Jon Marshall

"The DN bit is a check that, usually, PE routers use to check whether to install certain types of LSAs into a VRF and is used as a loop prevention method.

If your CE router is not running VRFs but using OSPF to connect to the PE router then you do not need that command anywhere.

If however you configure VRFs on your CE router then it now uses the same checks as the PE routers because it believes it is directly connected to the MPLS network in the way the PE is, even though it isn't.

And then you would need to use that command on your CE router.

So, put simply, you only need to use that command if your CE router is using "VRF-Lite" and OSPF is in use between the CE and PE routers.
..."

upvoted 2 times

Budh 1 year, 10 months ago

Selected Answer: D

Answer is D

upvoted 1 times

wts 1 year, 10 months ago

Selected Answer: D

capability vrf-lite command should be enabled:
- only on the CE router
- only when you have VRFs on your CE router

upvoted 4 times

error_909 2 years, 3 months ago

The given answer is correct D

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct
https://community.cisco.com/t5/routing/where-to-configure-the-quot-capability-vrf-lite-quot-on-ce-or-pe/td-p/2812305

upvoted 3 times

Masashi_O 2 years, 6 months ago

A is the answer, I think.

upvoted 1 times

Question #109 *Topic 1*

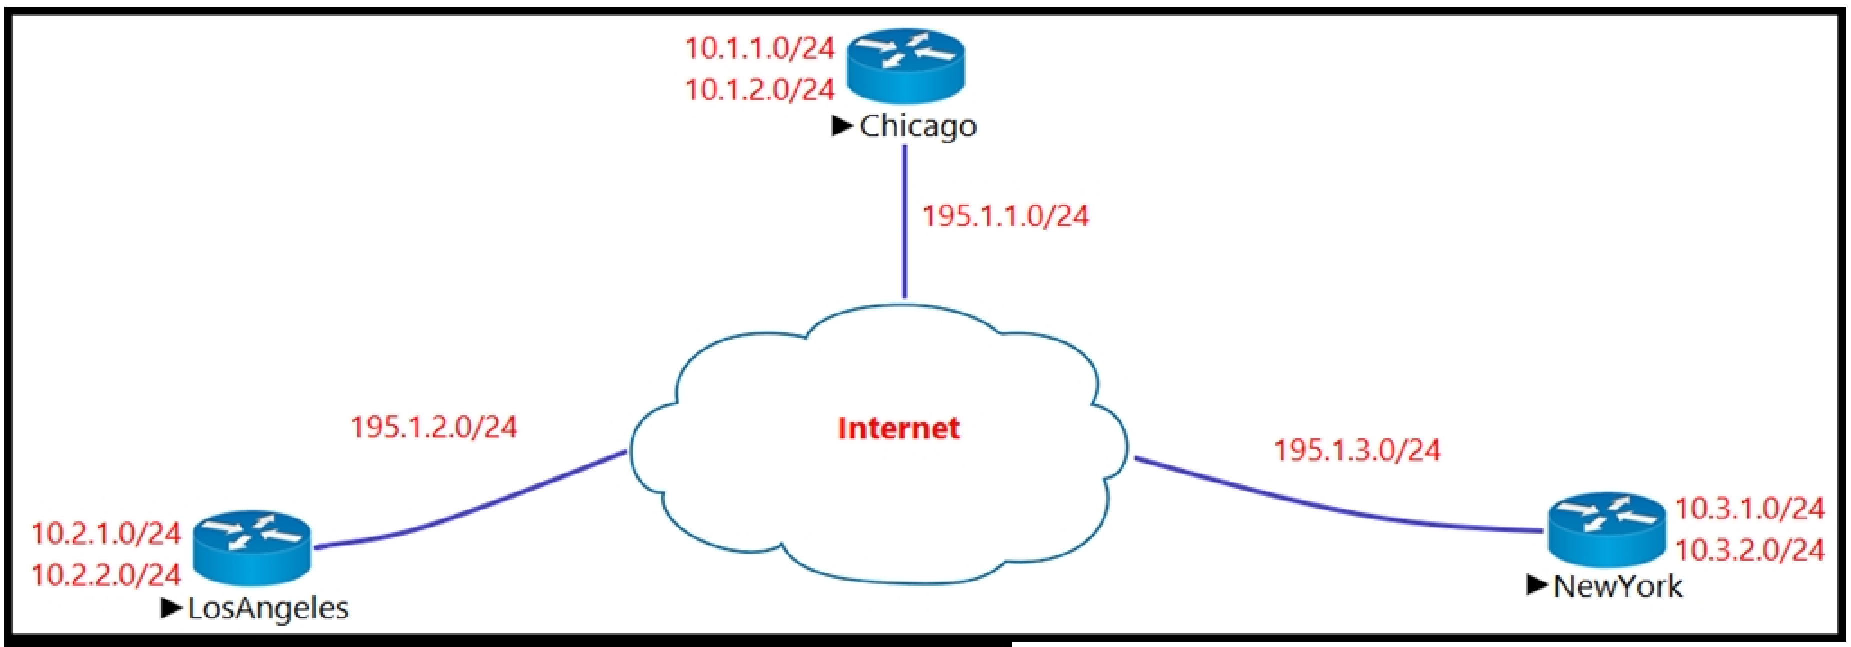

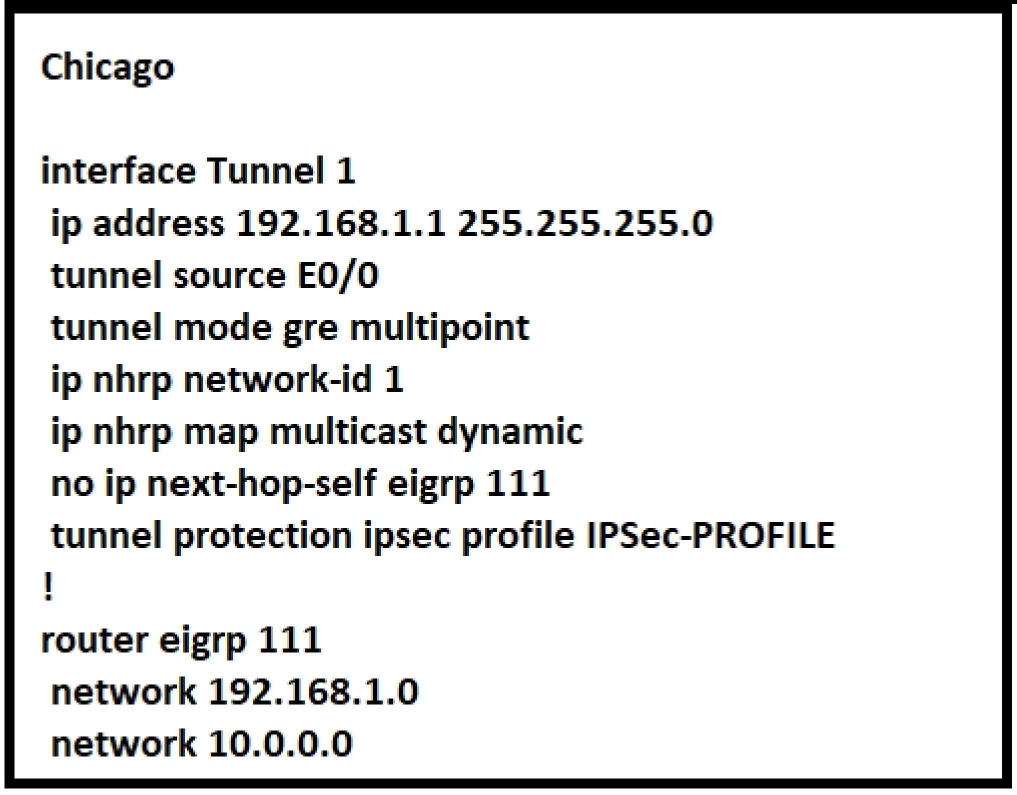

Chicago

```
interface Tunnel 1
 ip address 192.168.1.1 255.255.255.0
 tunnel source E0/0
 tunnel mode gre multipoint
 ip nhrp network-id 1
 ip nhrp map multicast dynamic
 no ip next-hop-self eigrp 111
 tunnel protection ipsec profile IPSec-PROFILE
!
router eigrp 111
 network 192.168.1.0
 network 10.0.0.0
```

Refer to the exhibit. The Los Angeles and New York routers are receiving routers from Chicago but not from each other.
Which configuration fixes the issue?

A. interface Tunnel1 no ip split-horizon eigrp 111

B. interface Tunnel1 ip next-hop-self eigrp 111

C. interface Tunnel1 tunnel mode ipsec ipv4

D. interface Tunnel1 tunnel protection ipsec profile IPSec-PROFILE

Correct Answer: *A*

Surfside92 Highly Voted 👍 2 years, 1 month ago

The given answer is correct - A

Its important here to work out that Chicago is the Hub router in a DMVWN network.
If Chicago was a spoke it would need a mapping to the hub - and that is not in the output - ie "ip nhrp map" command with relevant tunnel and WAN ip addresses ip addresses for the hub.

The hub over its tunnel1 interface learns the routes from LA - it wants to advertise the LA routes to New York - but those advertisements would be back out Tunnel1 - split horizon will not allow this. Split horizon does not allow advertising routes back out the interface a router received them on.
So the fix is to disable split horizon - a valid fix in certain scenarios.

Note this is a phase 2 DMVPN solution
In Phase 2, the Hub is still the “hub” for the control plane. All routes are learned through the hub. Spokes cannot exchange routes with each other directly.

A phase 3 DMVPN solution does not have this split horizon issue.

upvoted 13 times

Noproblem22 Most Recent ⊙ 1 year, 1 month ago

A is correct

upvoted 2 times

krn007 1 year, 11 months ago

Suggested Answer A is correct.
@Surfside92: Hi I checked DMVPN-Ph3 topology in gns3, even here "no ip split-horizon eigrp <as>" config is required on the HUB site tunnel interface.

upvoted 2 times

GReddy2323 2 years, 2 months ago

Can anybody give an explanation why A is the answer?

upvoted 1 times

cyrus777 1 year, 8 months ago

n this topology, Chicago router (Hub) will receive advertisements from Los Angeles (Spoke1) router on its tunnel interface. The problem here is that it also has a connection with New York (Spoke2) on that same tunnel interface. If we don't disable EIGRP split-horizon, then the Hub will not relay routes from Spoke1 to Spoke2 and the other way around. That is because it received those routes on interface Tunnel1 and therefore it cannot advertise back out that same interface (splithorizon rule). Therefore we must disable split-horizon on the Hub router to make sure the Spokes know about each other.

upvoted 7 times

Question #110

Topic 1

DRAG DROP -

Drag and drop the MPLS VPN device types from the left onto the definitions on the right.

Select and Place:

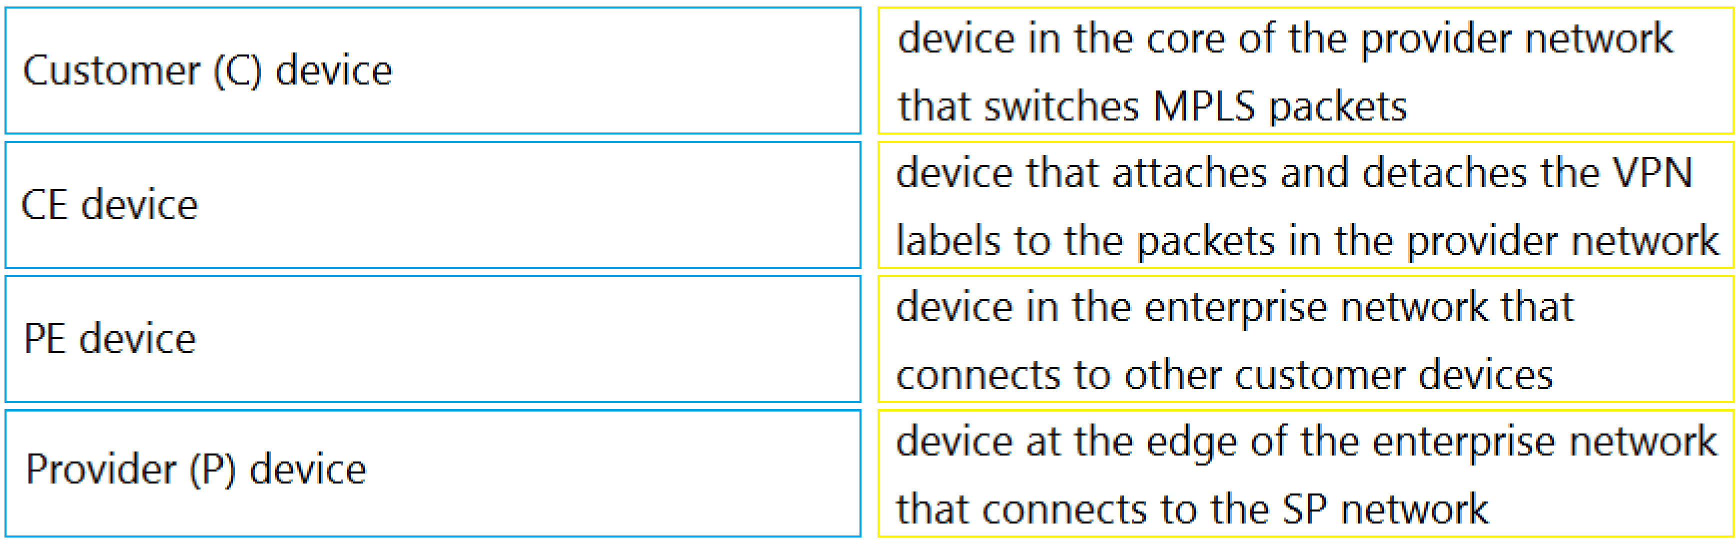

Customer (C) device	device in the core of the provider network that switches MPLS packets
CE device	device that attaches and detaches the VPN labels to the packets in the provider network
PE device	device in the enterprise network that connects to other customer devices
Provider (P) device	device at the edge of the enterprise network that connects to the SP network

Correct Answer:

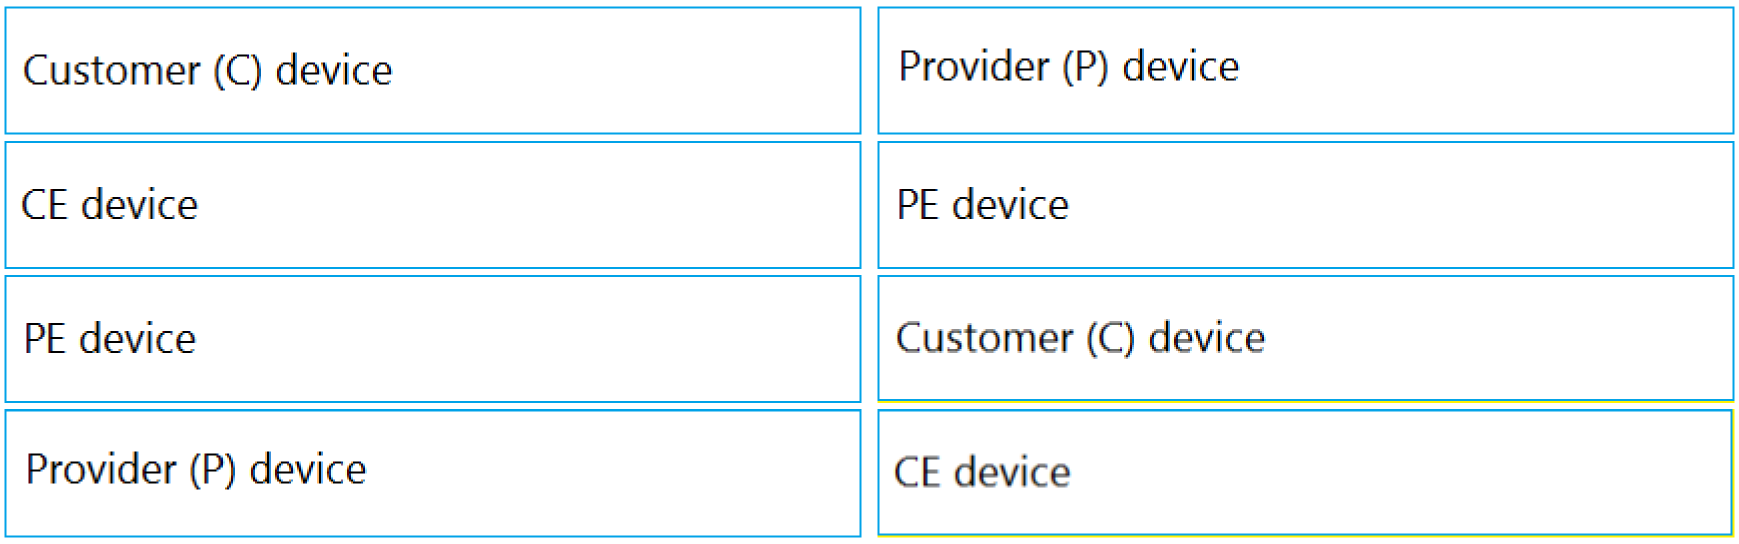

Customer (C) device	Provider (P) device
CE device	PE device
PE device	Customer (C) device
Provider (P) device	CE device

examShark Highly Voted 2 years, 4 months ago

The given answer is correct

upvoted 6 times

Question #111 Topic 1

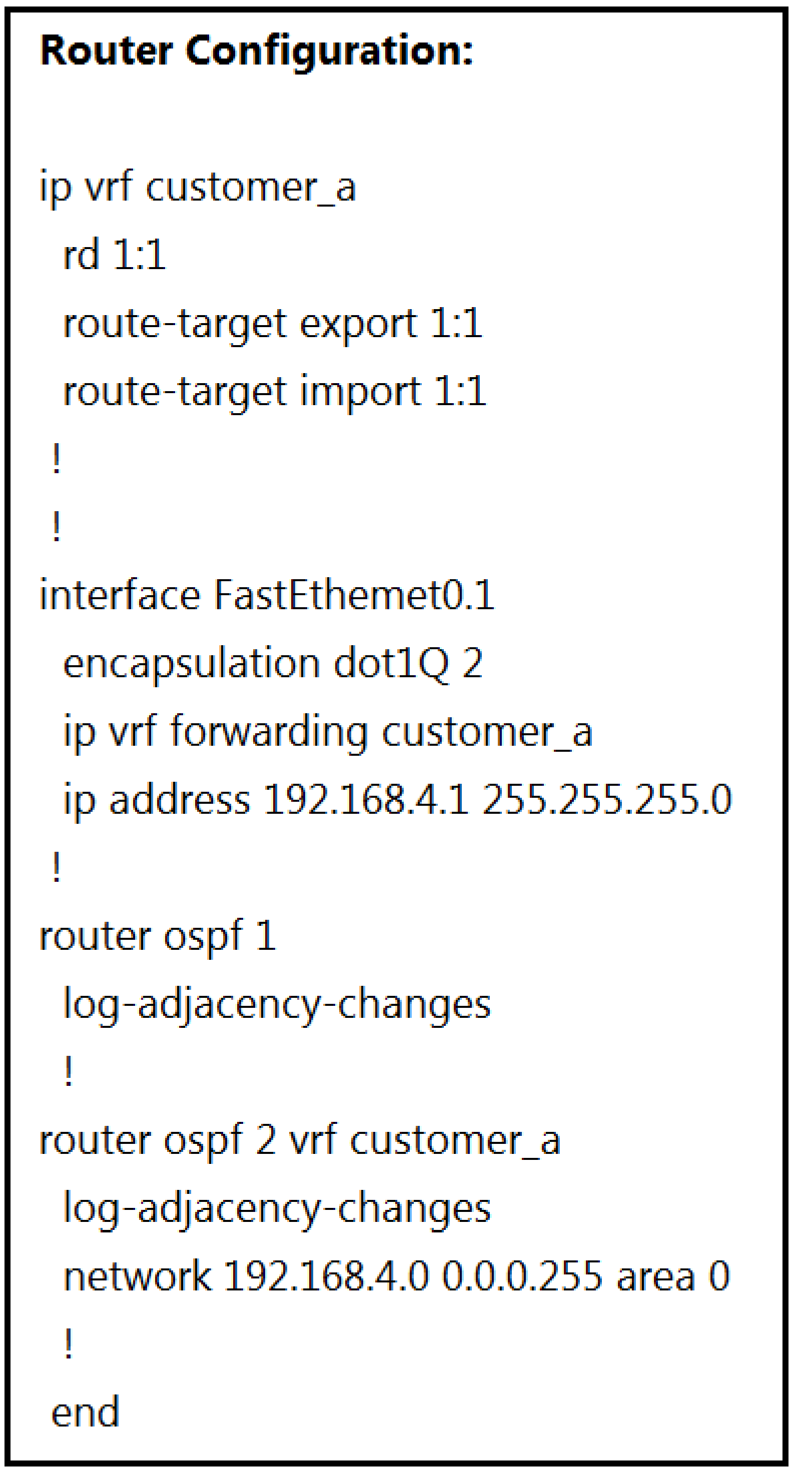

Router Configuration:

```
ip vrf customer_a
  rd 1:1
  route-target export 1:1
  route-target import 1:1
 !
 !
interface FastEthemet0.1
  encapsulation dot1Q 2
  ip vrf forwarding customer_a
  ip address 192.168.4.1 255.255.255.0
 !
router ospf 1
  log-adjacency-changes
  !
router ospf 2 vrf customer_a
  log-adjacency-changes
  network 192.168.4.0 0.0.0.255 area 0
  !
 end
```

Refer to the exhibit. The network administrator configured VRF lite for customer A. The technician at the remote site misconfigured VRF on the router.

Which configuration will resolve connectivity for both sites of customer_a?

A.

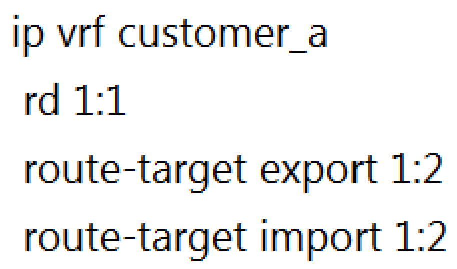

```
ip vrf customer_a
 rd 1:1
 route-target export 1:2
 route-target import 1:2
```

B.

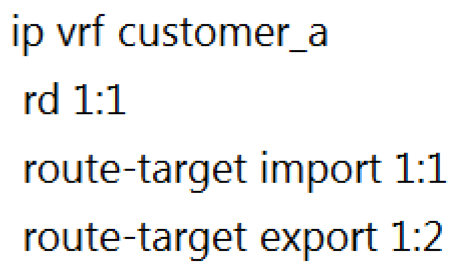

```
ip vrf customer_a
 rd 1:1
 route-target import 1:1
 route-target export 1:2
```

C.

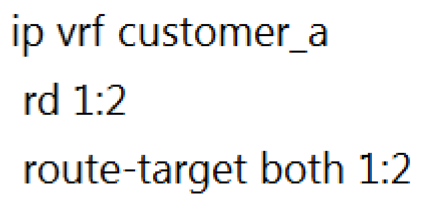

```
ip vrf customer_a
 rd 1:2
 route-target both 1:2
```

D.

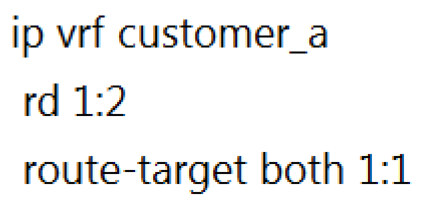

```
ip vrf customer_a
 rd 1:2
 route-target both 1:1
```

Correct Answer: *D*

⊟ **WhatNot** Highly Voted 2 years, 6 months ago

How do any of the answers have anything to do with the question ? Unless we see the import/export route target on the remote PE, any of these answers could be correct.

upvoted 9 times

⊟ **Pietjeplukgeluk** Most Recent 3 weeks, 6 days ago

For basic VRF-lite there is no need to specify RD or export import targets. So the questions missing context. It seems they do some kind of route leaking and require this at both ends. Just reading the question again, you should assume there is a requirement of route leaking with MP-BGP. Note sure, but personally this far more complex than basic VRF-lite.

upvoted 1 times

⊟ **inteldarvid** 5 months, 2 weeks ago

is D:
The network administrator configured VRF lite for customer A. The technician at the remote site misconfigured VRF on the router.
Which configuration will resolve connectivity for both sites of customer_a?

upvoted 1 times

⊟ **Caledonia** 1 year, 3 months ago

Without seeing the other side the router configs, it is impossible to decide what should be configured on CE router

upvoted 2 times

⊟ **Budh** 1 year, 10 months ago

Rt can be same on both routers, correct answet

upvoted 2 times

⊟ **examShark** 2 years, 4 months ago

The given answer is correct
rd local significance
rt same both ends

upvoted 4 times

⊟ **Masashi_O** 2 years, 6 months ago

A or D is the answer, but it is unclear whether this Config is on the network administrator's side or the remote technician's side.

upvoted 3 times

⊟ **Masashi_O** 2 years, 6 months ago

Since VRF customer_a is exporting and importing with a Route Target of 1:1, the remote device must also be exporting and importing with a Route Target of 1:1.
So, the answer is D.

upvoted 6 times

Question #112 *Topic 1*

What does the PE router convert the IPv4 prefix to within an MPLS VPN?

A. eBGP path association between the PE and CE sessions

B. prefix that combines the ASN, PE router-id, and IP prefix

C. 48-bit route combining the IP and PE router-id

D. VPN-IPv4 prefix combined with the 64-bit route distinguisher

Correct Answer: *D*

⊟ **Edwinmolinab** Highly Voted 1 year, 4 months ago

Explanation

The IP prefix is a member of the IPv4 address family. After the PE device learns the IP prefix, the PE converts it into a VPN-IPv4 prefix by combining it with an 8-byte route distinguisher (RD). The generated prefix is a member of the VPN-IPv4 address family. It uniquely identifies the customer address, even if the customer site is using globally nonunique (unregistered private) IP addresses.

upvoted 5 times

⊟ **Noproblem22** Most Recent 1 year, 1 month ago

D is correct

upvoted 2 times

⊟ **error_909** 2 years, 3 months ago

The given answer is correct D

upvoted 1 times

⊟ **examShark** 2 years, 4 months ago

The given answer is correct

upvoted 3 times

Question #113 Topic 1

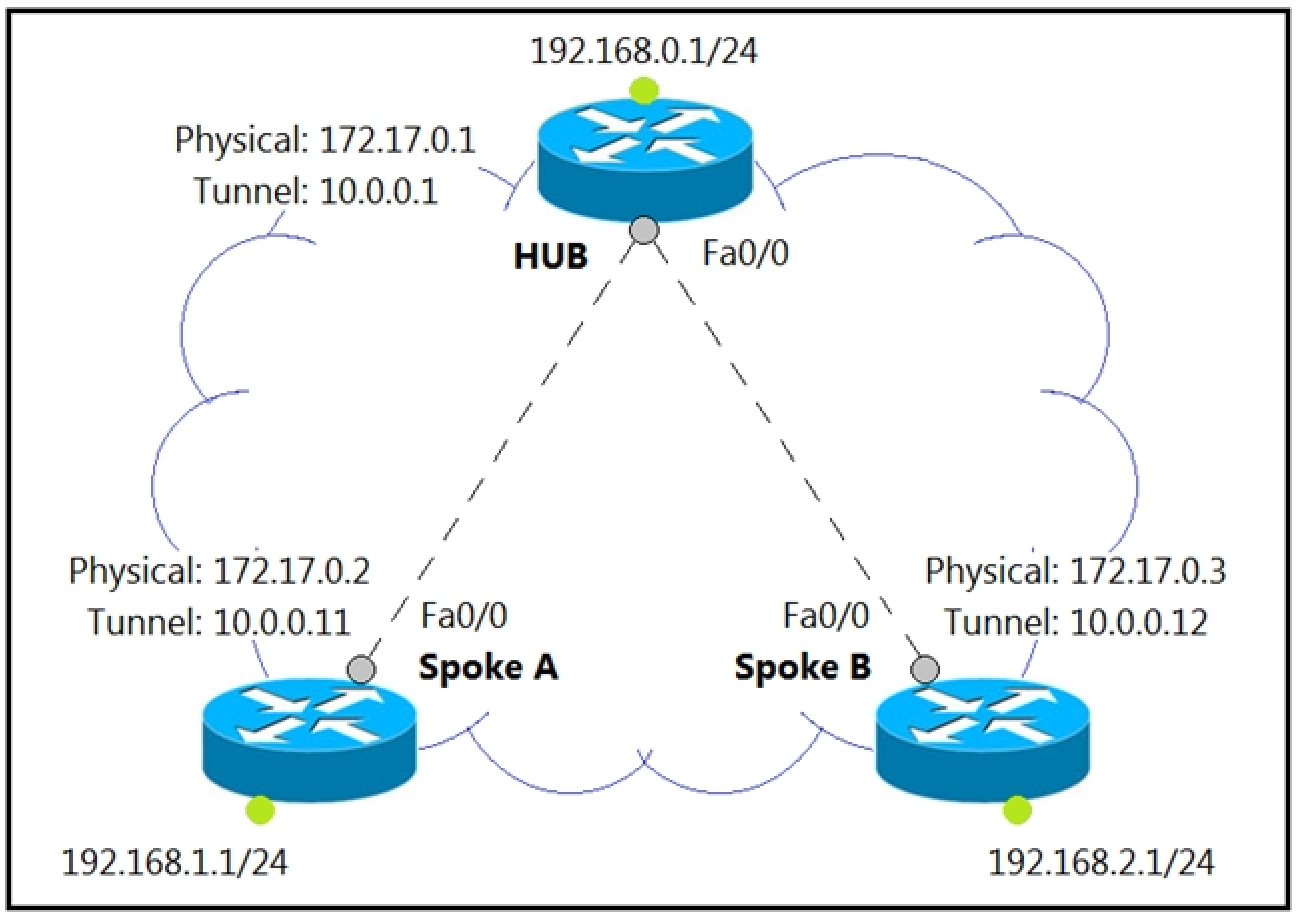

Refer to the exhibit. Which interface configuration must be configured on the HUB router to enable MVPN with mGRE mode?

A. interface Tunnel0 description mGRE - DMVPN Tunnel ip address 10.1.0.1 255.255.255.0 ip nhrp map multicast dynamic ip nhrp network-id 1 tunnel source 172.17.0.1 ip nhrp map 10.0.0.11 172.17.0.2 ip nhrp map 10.0.0.12 172.17.0.3 tunnel mode gre

B. interface Tunnel0 description mGRE - DMVPN Tunnel ip address 10.0.0.1 255.255.255.0 ip nhrp map multicast dynamic ip nhrp network-id 1 tunnel source 10.0.0.1 tunnel mode gre multipoint

C. interface Tunnel0 description mGRE - DMVPN Tunnel ip address 10.0.0.1 255.255.255.0 ip nhrp network-id 1 tunnel source 172.17.0.1 tunnel mode gre multipoint

D. interface Tunnel0 description mGRE - DMVPN Tunnel ip address 10.0.0.1 255.255.255.0 ip nhrp map multicast dynamic ip nhrp network-id 1 tunnel source 10.0.0.1 tunnel destination 172.17.0.2 tunnel mode gre multipoint

Correct Answer: *B*

Community vote distribution

C (100%)

DaanB Highly Voted 2 years, 8 months ago

Tunnel source IP can NOT be the IP address of the tunnel interface. The tunnel source IP should be, in this case, the IP address of the WAN interface.

upvoted 12 times

Brand Most Recent 3 months, 4 weeks ago

```
A.
interface Tunnel0 description mGRE - DMVPN Tunnel
ip address 10.1.0.1 255.255.255.0
ip nhrp map multicast dynamic
ip nhrp network-id 1
tunnel source 172.17.0.1
ip nhrp map 10.0.0.11 172.17.0.2
ip nhrp map 10.0.0.12 172.17.0.3
tunnel mode gre
B.
interface Tunnel0 description mGRE - DMVPN Tunnel
ip address 10.0.0.1 255.255.255.0
ip nhrp map multicast dynamic
ip nhrp network-id 1
tunnel source 10.0.0.1
tunnel mode gre multipoint
C.
interface Tunnel0 description mGRE - DMVPN Tunnel
ip address 10.0.0.1 255.255.255.0
ip nhrp network-id 1
tunnel source 172.17.0.1
tunnel mode gre multipoint
```

D.

```
interface Tunnel0 description mGRE - DMVPN Tunnel
ip address 10.0.0.1 255.255.255.0
ip nhrp map multicast dynamic
ip nhrp network-id 1
tunnel source 10.0.0.1
tunnel destination 172.17.0.2
tunnel mode gre multipoint
```

upvoted 2 times

Chiaretta 5 months, 1 week ago

Selected Answer: C

C is correct

upvoted 1 times

forccnp 9 months, 3 weeks ago

Selected Answer: C

C is the correct answer

upvoted 2 times

Koume 11 months ago

Selected Answer: C

Even is missing ip nhrp multicast dynamic. Seems the most correct as all command are valir for HUB

upvoted 3 times

Pietjeplukgeluk 3 weeks, 6 days ago

I agree, it C should have "ip nhrp map multicast dynamic" added to be fully correct. Anyway it seems better than B as that has the wrong source ip specified.

upvoted 1 times

ChillingAgain 1 year, 1 month ago

Selected Answer: C

C is correct

upvoted 1 times

JOKERR 1 year, 6 months ago

Selected Answer: C

C is correct.

upvoted 2 times

Hack4 1 year, 10 months ago

C is correct

upvoted 1 times

Budh 1 year, 10 months ago

Selected Answer: C

C is correct

upvoted 1 times

tyh391 1 year, 10 months ago

Selected Answer: C

As in Discussion

upvoted 1 times

[Removed] 1 year, 10 months ago

Selected Answer: C

Answer is C. The question specifies how to enable MGRE Mode. Now your tunnel source cannot be a tunnel IP. It can either be a physical interface or a physical interface IP

upvoted 2 times

Carl1999 1 year, 10 months ago

Selected Answer: C

C is correct not B.
ip address 10.0.0.1 255.255.255.0 ->tunnel ip address
tunnel source 172.17.0.1 ->physical ip address

upvoted 3 times

wts 1 year, 10 months ago

A - invalid tunnel ip-address.
B - invalid tunnel source ip-address.
C - multicast not enabled.
D - invalid tunnel source ip-address.

Apparently, everyone is choosing between the wrong ip-address and the missing multicast enable command.

upvoted 4 times

TOKI07 1 year, 11 months ago

C is correct.
"ip nhrp map multicast dynamic" needs to be configured mainly when dynamic routing protocols using multicast is used.

upvoted 3 times

gndrx78 1 year, 11 months ago

What about "tunnel source 172.17.0.1"? Seems wrong.

upvoted 1 times

gndrx78 1 year, 11 months ago

My mistake, tunnel source must be public IP address so only A and C answers could be the right one. But C is the only answers with proper command "tunnel mode gre multipoint", so it is C

upvoted 3 times

C_Tw21 2 years, 1 month ago

C is correct.
B is not correct as it uses the wrong tunnel-source. Should be the WAN-Interface , not the tunnel-ip itself.
The "ip nhrp map multicast dynamic" is used on the spokes , and should point toi the hub.

upvoted 3 times

Raider1 2 years, 2 months ago

B is answer. You need ip nhrp map multicast dynamic on the HUB

upvoted 2 times

OakA1 2 years, 2 months ago

B is nonsense. C is correct

upvoted 1 times

Question #114 *Topic 1*

How are MPLS Layer 3 VPN services deployed?

A. The RD and RT values must match under the VRF.

B. The import and export RT values under a VRF must always be the same.

C. The label switch path must be available between the local and remote PE routers.

D. The RD and RT values under a VRF must match on the remote PE router.

Correct Answer: *C*

Community vote distribution

C (100%)

ysue 5 months, 2 weeks ago

Answer is correct.

https://www.cisco.com/c/en/us/td/docs/dcn/nx-os/nexus9000/103x/configuration/label-switching/cisco-nexus-9000-series-nx-os-label-switching-configuration-guide-103x/m-configuring-mpls-layer-3-vpns.pdf

upvoted 1 times

juliop 11 months, 1 week ago

Why B is not correct?

upvoted 1 times

palihaff 1 year, 11 months ago

Selected Answer: C

The given answer is correct

upvoted 2 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #115 *Topic 1*

Which IGPs are supported by the MPLS LDP autoconfiguration feature?

A. IS-IS and RIPv2

B. RIPv2 and OSPF

C. OSPF and EIGRP

D. OSPF and IS-IS

Correct Answer: *D*

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/15-s/mp-ldp-15-s-book/mp-ldp-autoconfig.pdf

SnoopDD 2 months ago

The MPLS LDP Autoconfiguration feature enables you to globally enable Label Distribution Protocol (LDP) on every interface associated with an Interior Gateway Protocol (IGP) instance. This feature is supported on Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) IGPs. It provides a means to block LDP from being enabled on interfaces that you do not want enabled. The goal of the MPLS LDP Autoconfiguration feature is to make configuration easier, faster, and error free.

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 3 times

Refer to the exhibit.

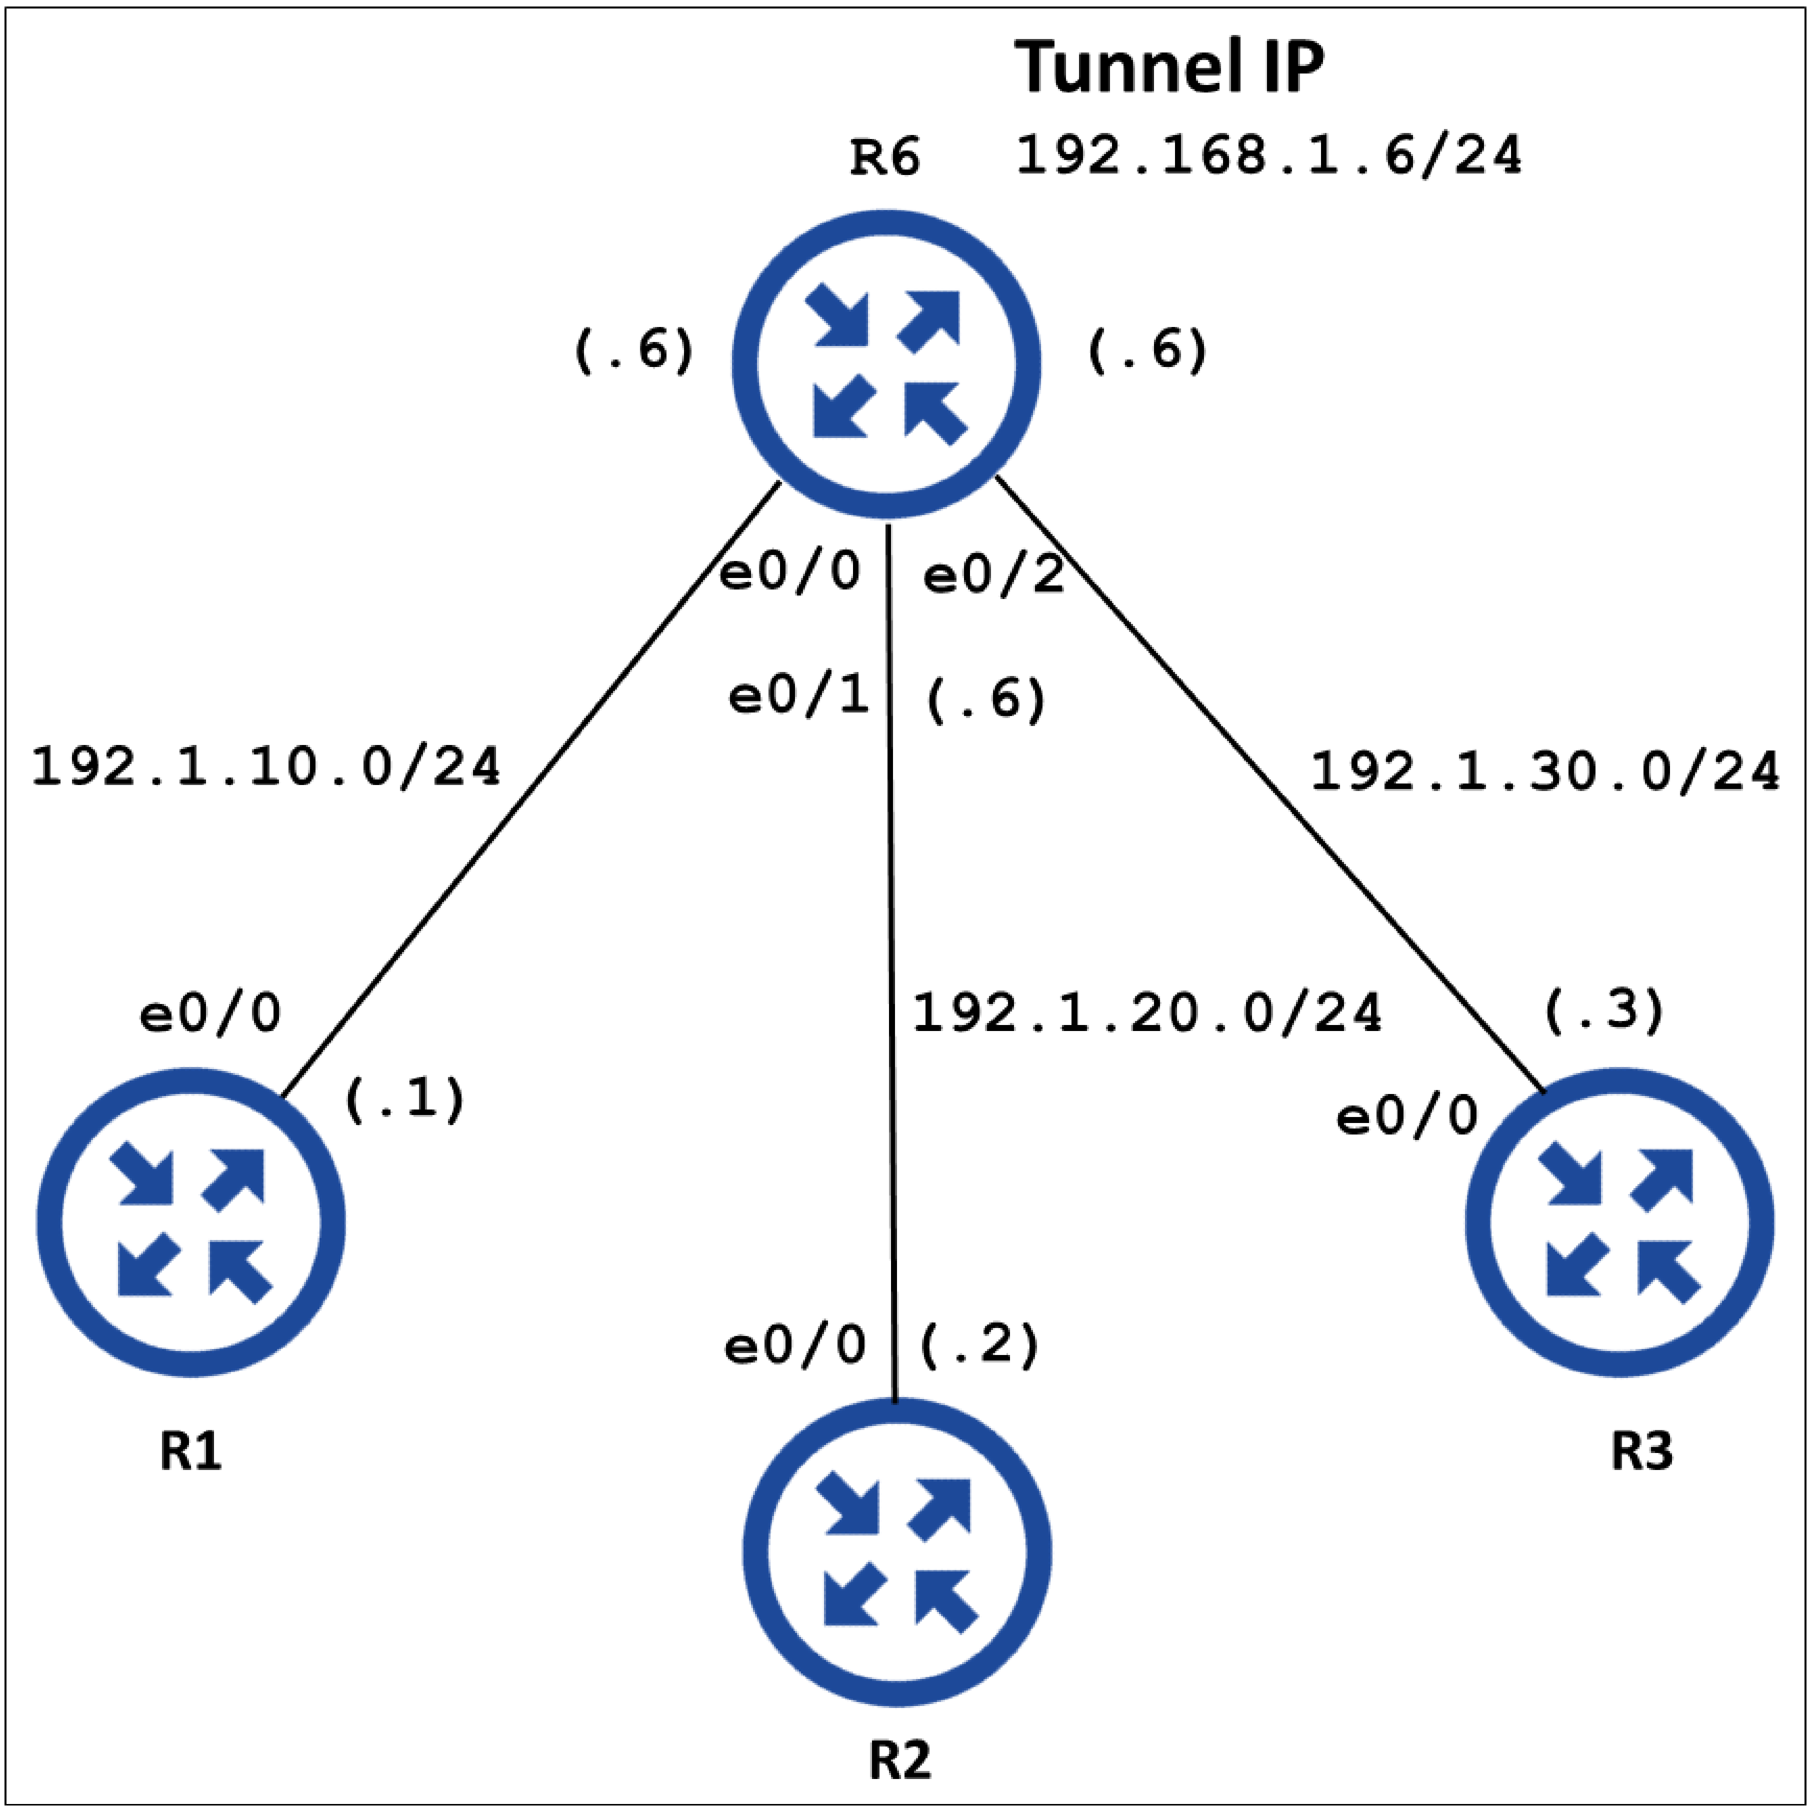

An engineer must establish multipoint GRE tunnels between hub router R6 and branch routers R1, R2, and R3.
Which configuration accomplishes this task on R1?

A. interface Tunnel 1 ip address 192.168.1.1 255.255.255.0 tunnel source e0/0 tunnel mode gre multipoint ip nhrp nhs 192.168.1.6 ip nhrp map 192.168.1.6 192.1.10.1 ip nhrp map 192.168.1.2 192.1.20.2 ip nhrp map 192.168.1.3 192.1.30.3

B. interface Tunnel 1 ip address 192.168.1.1 255.255.255.0 tunnel source e0/1 tunnel mode gre multipoint ip nhrp nhs 192.168.1.6 ip nhrp map 192. 168.1.6 192.1.10.6

C. interface Tunnel 1 ip address 192.168.1.1 255.255.255.0 tunnel source e0/0 tunnel mode gre multipoint ip nhrp network-id 1 ip nhrp nhs 192.168.1.6 ip nhrp map 192.168.1.6 192.1.10.6

D. interface Tunnel 1 ip address 192.168.1.1 255. 255.255.0 tunnel source e0/1 tunnel mode gre multipoint ip nhrp network-id 1 ip nhrp nhs 192.168.1.6 ip nhrp map 192.168.1.6 192.1.10.1 ip nhrp map 192.168.1.2 192.1.20.2 ip nhrp map 192.168.1.3 192.1.30.3

Correct Answer: *C*

Community vote distribution

C (100%)

Brand 3 months, 4 weeks ago

Selected Answer: C

A.

```
interface Tunnel 1
ip address 192.168.1.1 255.255.255.0
tunnel source e0/0
tunnel mode gre multipoint
ip nhrp nhs 192.168.1.6
ip nhrp map 192.168.1.6 192.1.10.1
ip nhrp map 192.168.1.2 192.1.20.2
ip nhrp map 192.168.1.3 192.1.30.3
```

B.

```
interface Tunnel 1
ip address 192.168.1.1 255.255.255.0
tunnel source e0/1
tunnel mode gre multipoint
ip nhrp nhs 192.168.1.6
ip nhrp map 192. 168.1.6 192.1.10.6
```

C.

```
interface Tunnel 1
ip address 192.168.1.1 255.255.255.0
tunnel source e0/0
tunnel mode gre multipoint
ip nhrp network-id 1
ip nhrp nhs 192.168.1.6
ip nhrp map 192.168.1.6 192.1.10.6
```

D.

```
interface Tunnel 1
ip address 192.168.1.1 255. 255.255.0
tunnel source e0/1
tunnel mode gre multipoint
ip nhrp network-id 1
ip nhrp nhs 192.168.1.6
ip nhrp map 192.168.1.6 192.1.10.1
ip nhrp map 192.168.1.2 192.1.20.2
ip nhrp map 192.168.1.3 192.1.30.3
```

upvoted 1 times

Carl1999 1 year, 10 months ago

C is correct
B doesn't have a network iD command.

upvoted 2 times

VVdouble 1 year, 10 months ago

B is not wrong because of the missing network ID command, but because it uses e0/1 as tunnel source on R1 and it should be e0/0

upvoted 8 times

palihaff 1 year, 11 months ago

C is correct

upvoted 2 times

Question #117 *Topic 1*

How is VPN routing information distributed in an MPLS network?

A. The top level of the customer data packet directs it to the correct CE device.

B. It is established using VPN IPsec peers.

C. It is controlled through the use of RD.

D. It is controlled using of VPN target communities.

Correct Answer: *D*

Reference:

https://www.ccexpert.us/mpls-design/chapter-5-packetbased-mpls-vpns.html

MrThinMints Highly Voted 1 year, 11 months ago

Provided answer is correct.

upvoted 7 times

ZamanR Most Recent 5 days, 20 hours ago

D is the answer

The distribution of virtual private network (VPN) routing information is controlled through the use of VPN route target communities, implemented by Border Gateway Protocol (BGP) extended communities.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/15-mt/mp-l3-vpns-15-mt-book/mp

upvoted 1 times

Cyril_the_Squirl 4 months, 1 week ago

C is correct

upvoted 2 times

Question #118

Topic 1

IPv6 is enabled in the infrastructure to support customers with an IPv6 network over WAN and to connect the head office to branch offices in the local network.
One of the customers is already running IPv6 and wants to enable IPv6 over the DMVPN network infrastructure between the headend and branch sites.
Which configuration command must be applied to establish an mGRE IPv6 tunnel neighborship?

A. ipv6 nhrp holdtime 30

B. tunnel mode gre multipoint ipv6

C. ipv6 unicast-routing

D. tunnel protection mode ipv6

Correct Answer: *B*
Reference:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/interface/configuration/xe-16/ir-xe-16-book/ip6-mgre-tunls.pdf

Community vote distribution

B (100%)

⊟ **inteldarvid** 5 months, 2 weeks ago

Selected Answer: B

option B is correct:

https://www.pearsonitcertification.com/articles/article.aspx?p=3129283&seqNum=6

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/interface/configuration/xe-3s/ir-xe-3s-book/ir-gre-ipv6-tunls-xe.pdf

upvoted 2 times

Question #119

Topic 1

What is a characteristic of Layer 3 MPLS VPNs?

A. Traffic engineering capabilities provide QoS and SLAs.

B. Traffic engineering supports multiple IGP instances.

C. LSP signaling requires the use of unnumbered IP links for traffic engineering.

D. Authentication is performed by using digital certificates or preshared keys.

Correct Answer: *A*
Reference:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/15-mt/mp-te-diffserv-15-mt-book/mp-te-diffserv-aw.html

Community vote distribution

A (100%)

⊟ **inteldarvid** 5 months, 2 weeks ago

Selected Answer: A

Option A corerct:

https://vceguide.com/what-is-a-characteristic-of-layer-3-mpls-vpns/
https://itexamanswers.net/question/what-is-a-characteristic-of-layer-3-mpls-vpns

upvoted 1 times

Question #120 *Topic 1*

How does an MPLS Layer 3 VPN differentiate the IP address space used between each VPN?

A. by RT

B. by address family

C. by RD

D. by MP-BGP

Correct Answer: *C*

Community vote distribution

C (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: C

YES, C CORRECT

upvoted 1 times

PimplePooper 12 months ago

Selected Answer: C

C s correct.

upvoted 2 times

Question #121 *Topic 1*

Which OSI model is used to insert an MPLS label?

A. between Layer 2 and Layer 3

B. between Layer 5 and Layer 6

C. between Layer 1 and Layer 2

D. between Layer 3 and Layer 4

Correct Answer: *A*

Community vote distribution

A (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: A

yes, OPTION A layer 2.5

layer 1
layer2
----->HERE MPLS 2.5 <------ because mpls is fasther control plane L3
layer3
layer4
layer5
layer 6
layer7

upvoted 2 times

GreatDane 1 year, 5 months ago

Ref: Multiprotocol Label Switching – Wikipedia

"...
Role and functioning

MPLS operates at a layer that is generally considered to lie between traditional definitions of OSI Layer 2 (data link layer) and Layer 3 (network layer), and thus is often referred to as a layer 2.5 protocol.
..."

A. between Layer 2 and Layer 3

Correct answer.

B. between Layer 5 and Layer 6

Wrong answer.

C. between Layer 1 and Layer 2

Wrong answer.

D. between Layer 3 and Layer 4

Wrong answer.

upvoted 2 times

Question #122 *Topic 1*

Which function does LDP provide in an MPLS topology?

A. It enables a MPLS topology to connect multiple VPNs to P routers.

B. It provides hop-by-hop forwarding in an MPLS topology for LSRs.

C. It exchanges routes for MPLS VPNs across different VRFs.

D. It provides a means for LSRs to exchange IP routes.

Correct Answer: *B*

Community vote distribution

B (100%)

⊟ **inteldarvid** 5 months, 2 weeks ago

Selected Answer: B

correct B

upvoted 1 times

⊟ **Malasxd** 7 months, 1 week ago

Selected Answer: B

B is correct. The others does not make sense

upvoted 1 times

⊟ **GreatDane** 1 year, 5 months ago

Ref: Multiprotocol Label Switching (MPLS) on Cisco Routers

“...
Distribution of Label Bindings
...
• Label Distribution Protocol (LDP) - Enables peer LSRs in an MPLS network to exchange label binding information for supporting hop-by-hop forwarding in an MPLS network.
...”

A. It enables a MPLS topology to connect multiple VPNs to P routers.

Wrong answer.

B. It provides hop-by-hop forwarding in an MPLS topology for LSRs.

Correct answer.

C. It exchanges routes for MPLS VPNs across different VRFs.

Wrong answer.

D. It provides a means for LSRs to exchange IP routes.

Wrong answer.

upvoted 4 times

Which mechanism provides traffic segmentation within a DMVPN network?

A. BGP

B. IPsec

C. MPLS

D. RSVP

Correct Answer: *C*

Community vote distribution

inteldarvid 5 months, 2 weeks ago

Selected Answer: C

option correct is C:
Prerequisites for Dynamic Multipoint VPN (DMVPN)

To enable 2547oDMPVN--Traffic Segmentation Within DMVPN you must configure multiprotocol label switching (MPLS) by using the mpls ip command.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_dmvpn/configuration/15-mt/sec-conn-dmvpn-15-mt-book/sec-conn-dmvpn-dmvpn.html

upvoted 1 times

HungarianDish 7 months, 4 weeks ago

Selected Answer: C

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_dmvpn/configuration/15-mt/sec-conn-dmvpn-15-mt-book/sec-conn-dmvpn-dmvpn.html
"(To enable) Traffic Segmentation Within DMVPN you must configure multiprotocol label switching (MPLS) by using the mpls ip command."

upvoted 2 times

azzawim 9 months, 1 week ago

wrong its B

upvoted 2 times

jthompaf 1 year, 7 months ago

Given answer is correct:
https://www.cisco.com/c/en/us/td/docs/ios/ios_xe/sec_secure_connectivity/configuration/guide/convert/sec_dmvpn_xe_3s_book/sec_DMVPN_xe.html#:~:text=To%20use%20the%202547oDMPVN%E2%80%94Traffic%20Segmentation%20Within%20DMVPN%20feature%20you%20must%20configure%20Multiprotocol%20Label%20Switching%20(MPLS)%20by%20using%20the%20mpls%20ip%20command.

upvoted 3 times

Question #124

Topic 1

Refer to the exhibit. Which configuration denies Telnet traffic to router 2 from 198A:0:200C::1/64?

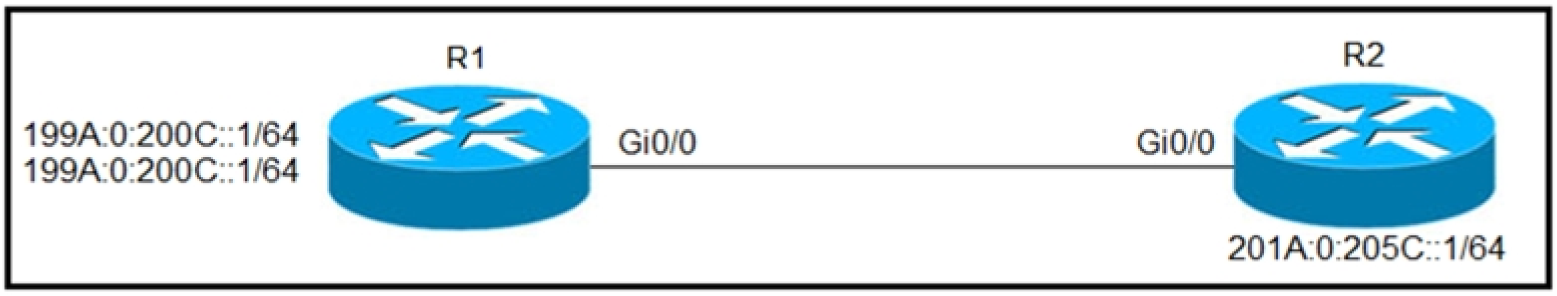

A.

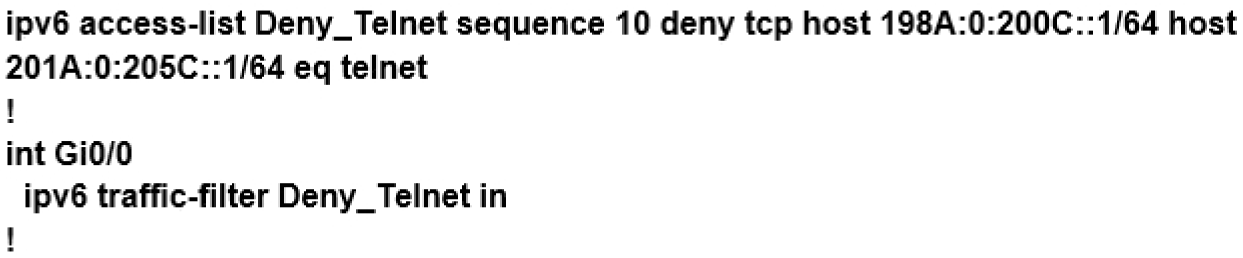

```
ipv6 access-list Deny_Telnet sequence 10 deny tcp host 198A:0:200C::1/64 host
201A:0:205C::1/64 eq telnet
!
int Gi0/0
  ipv6 traffic-filter Deny_Telnet in
!
```

B.

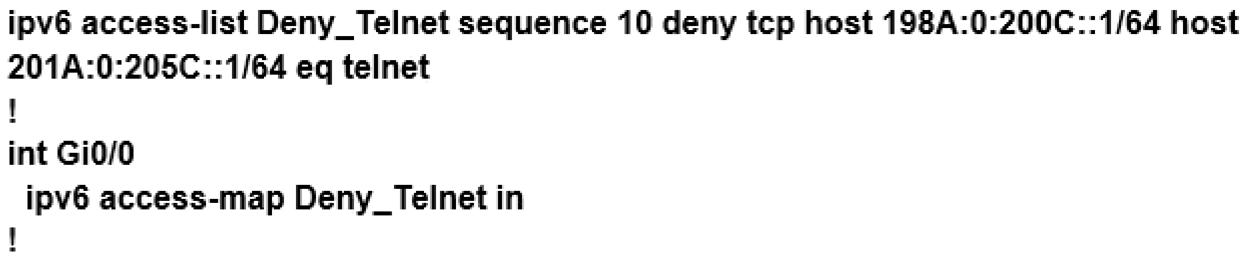

```
ipv6 access-list Deny_Telnet sequence 10 deny tcp host 198A:0:200C::1/64 host
201A:0:205C::1/64 eq telnet
!
int Gi0/0
  ipv6 access-map Deny_Telnet in
!
```

C.

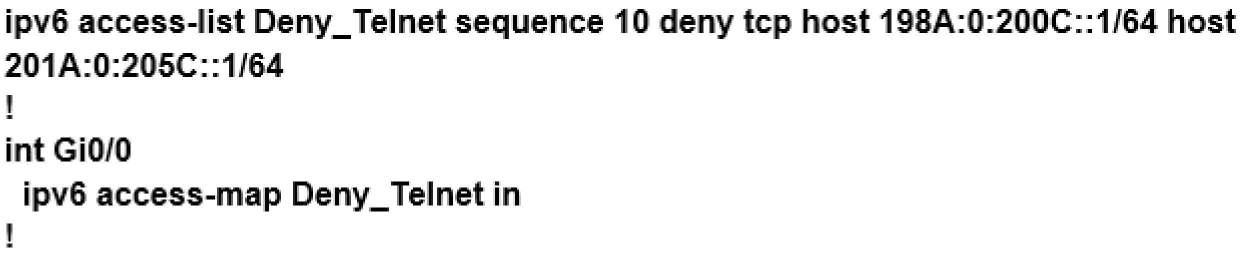

```
ipv6 access-list Deny_Telnet sequence 10 deny tcp host 198A:0:200C::1/64 host
201A:0:205C::1/64
!
int Gi0/0
  ipv6 access-map Deny_Telnet in
!
```

D.

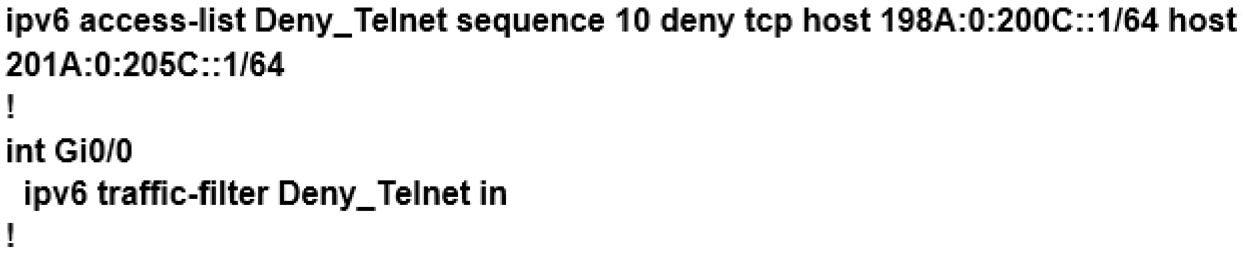

```
ipv6 access-list Deny_Telnet sequence 10 deny tcp host 198A:0:200C::1/64 host
201A:0:205C::1/64
!
int Gi0/0
  ipv6 traffic-filter Deny_Telnet in
!
```

Correct Answer: *A*

⊟ **[Removed]** 4 months, 1 week ago

A is the best answer, but incomplete.
B and C are using the wrong syntax to apply the ipv6 acl in the interface, along with C missing the telnet port in the destination portion
D is allowing all type of traffic from the indicated network.

upvoted 1 times

⊟ **inteldarvid** 5 months, 2 weeks ago

Option A

here this example from cisco.com
https://www.cisco.com/c/en/us/support/docs/ip/ip-version-6/113126-ipv6-acl-00.html

upvoted 1 times

⊟ **Dacusai** 7 months, 2 weeks ago

A is the correct answer, but still the Access list is missing another entry to permit the rest of the traffic. In this case all traffic will be denied due to the implicit deny at the end of the Access List.

upvoted 2 times

⊟ **Malasxd** 8 months ago

A seems more correct

upvoted 1 times

⊟ **YaPet** 1 year, 10 months ago

A is correct, because we need to deny only telnet traffic from R1, no any another traffic is mentioned in the question

upvoted 1 times

Carl1999 1 year, 10 months ago

198A:0:200C::1/64?
199A:0:200C::1/64?

upvoted 1 times

Carl1999 1 year, 10 months ago

and,,Is "permit ipv6 any any" unnecessary?
umm

upvoted 1 times

studybuddy10 2 years, 1 month ago

A is most correct, these ACLs still need a permit statement or they block all traffic. So D also works, B and C are bad syntax as they should be traffic-filter and not access-map.

upvoted 2 times

Carl1999 1 year, 10 months ago

D cannot communicate because it is denyed with implicit permission.

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #125 Topic 1

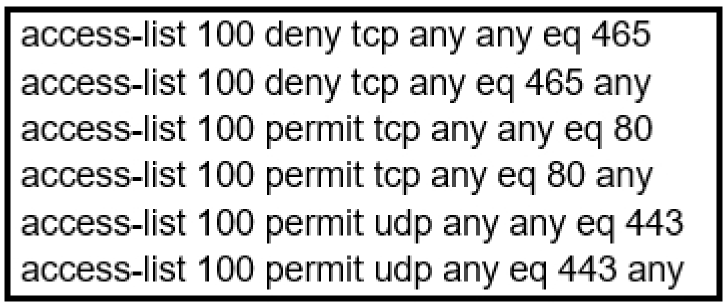

```
access-list 100 deny tcp any any eq 465
access-list 100 deny tcp any eq 465 any
access-list 100 permit tcp any any eq 80
access-list 100 permit tcp any eq 80 any
access-list 100 permit udp any any eq 443
access-list 100 permit udp any eq 443 any
```

Refer to the exhibit. During troubleshooting it was discovered that the device is not reachable using a secure web browser.
What is needed to fix the problem?

A. permit tcp port 443

B. permit udp port 465

C. permit tcp port 465

D. permit tcp port 22

Correct Answer: *A*

Community vote distribution

A (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: A

port 443 is HTTPS but work only TCP not UDP
Option A is correct

upvoted 2 times

Nhan 1 year, 6 months ago

https => port 443, http => port 80,

upvoted 2 times

AliMo123 2 years, 1 month ago

A is correct
port 443 uses TCP not UDP

upvoted 4 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

DRAG DROP -

Drag and drop the packet types from the left onto the correct descriptions on the right.

Select and Place:

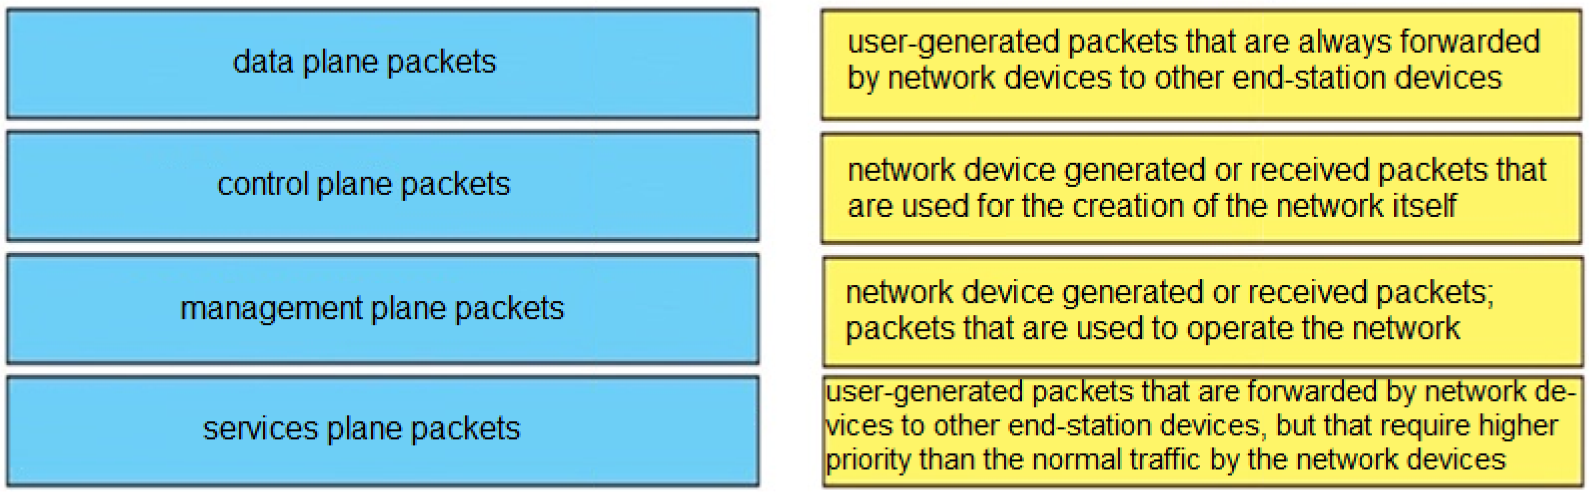

data plane packets

control plane packets

management plane packets

services plane packets

user-generated packets that are always forwarded by network devices to other end-station devices

network device generated or received packets that are used for the creation of the network itself

network device generated or received packets; packets that are used to operate the network

user-generated packets that are forwarded by network devices to other end-station devices, but that require higher priority than the normal traffic by the network devices

Correct Answer:

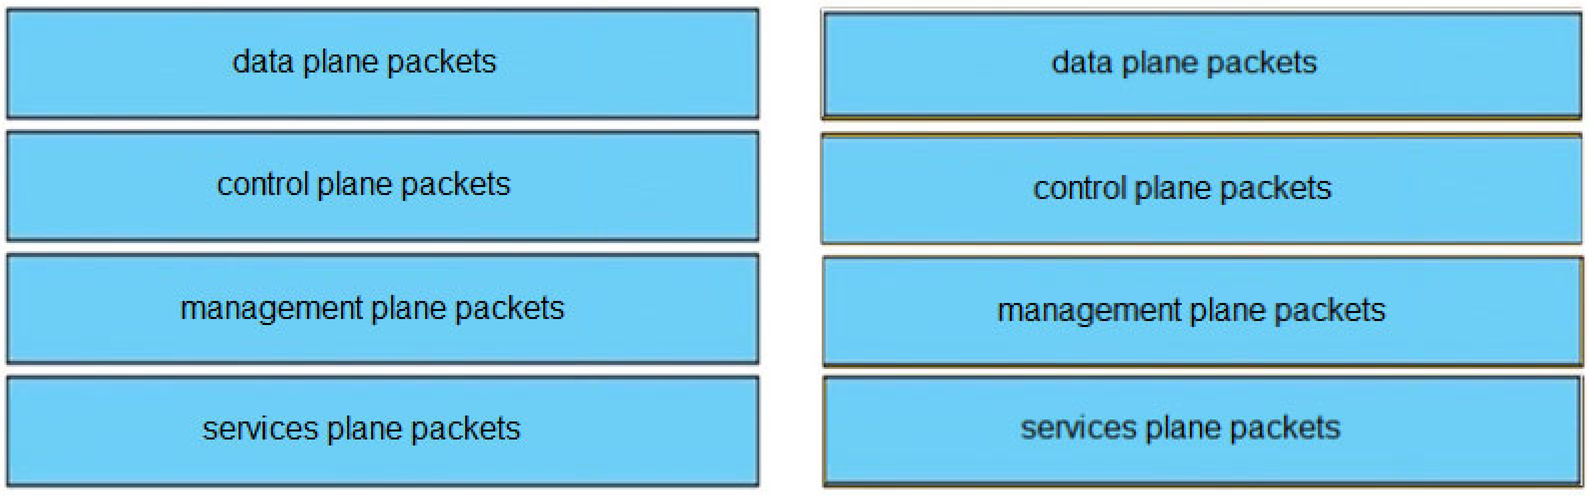

data plane packets

control plane packets

management plane packets

services plane packets

data plane packets

control plane packets

management plane packets

services plane packets

inteldarvid 5 months, 2 weeks ago

correct:

https://www.networktut.com/control-plane-policing-copp-tutorial

upvoted 1 times

guy276465281819372 5 months, 4 weeks ago

given answer is correct and description is clear.

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

ssbipa6 2 years, 8 months ago

they just go straight in their raws?

upvoted 3 times

Wesgo 2 years, 8 months ago

Yes, and the descriptions are very clear.

upvoted 2 times

Question #127 Topic 1

DRAG DROP -

Drag and drop the addresses from the left onto the correct IPv6 filter purposes on the right.

Select and Place:

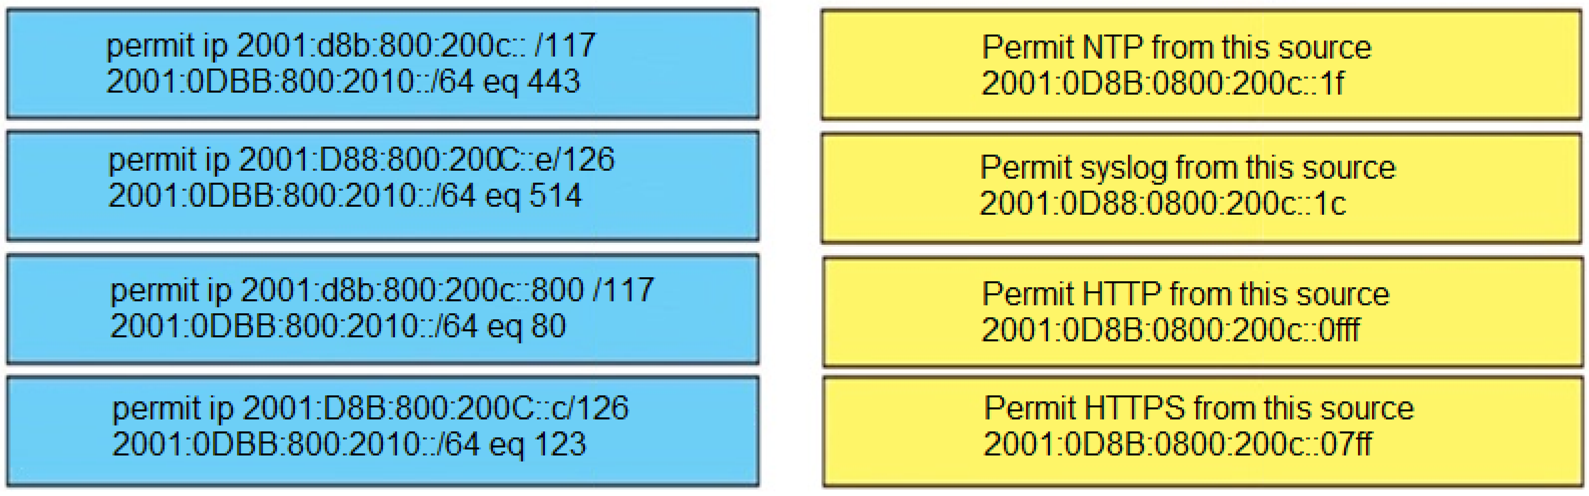

permit ip 2001:d8b:800:200c:: /117
2001:0DBB:800:2010::/64 eq 443

permit ip 2001:D88:800:200C::e/126
2001:0DBB:800:2010::/64 eq 514

permit ip 2001:d8b:800:200c::800 /117
2001:0DBB:800:2010::/64 eq 80

permit ip 2001:D8B:800:200C::c/126
2001:0DBB:800:2010::/64 eq 123

Permit NTP from this source
2001:0D8B:0800:200c::1f

Permit syslog from this source
2001:0D88:0800:200c::1c

Permit HTTP from this source
2001:0D8B:0800:200c::0fff

Permit HTTPS from this source
2001:0D8B:0800:200c::07ff

Correct Answer:

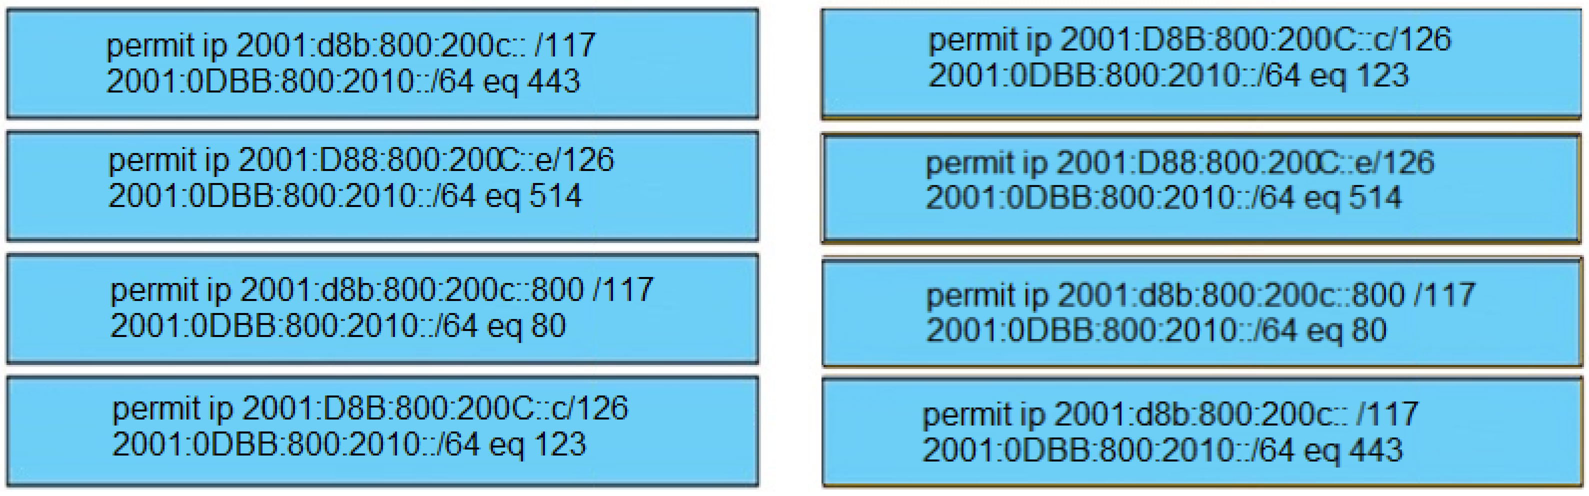

permit ip 2001:d8b:800:200c:: /117
2001:0DBB:800:2010::/64 eq 443

permit ip 2001:D88:800:200C::e/126
2001:0DBB:800:2010::/64 eq 514

permit ip 2001:d8b:800:200c::800 /117
2001:0DBB:800:2010::/64 eq 80

permit ip 2001:D8B:800:200C::c/126
2001:0DBB:800:2010::/64 eq 123

permit ip 2001:D8B:800:200C::c/126
2001:0DBB:800:2010::/64 eq 123

permit ip 2001:D88:800:200C::e/126
2001:0DBB:800:2010::/64 eq 514

permit ip 2001:d8b:800:200c::800 /117
2001:0DBB:800:2010::/64 eq 80

permit ip 2001:d8b:800:200c:: /117
2001:0DBB:800:2010::/64 eq 443

OhBee 1 year, 11 months ago

Wait...the masks are wrong for the NTP and syslog scenario.

Also, the "tcp" statement is missing on all of these.

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Jack1188 3 years, 4 months ago

this is correct answer

upvoted 3 times

Question #128 Topic 1

Refer to the exhibit. An engineer is trying to configure local authentication on the console line, but the device is trying to authenticate using TACACS+.

Which action produces the desired configuration?

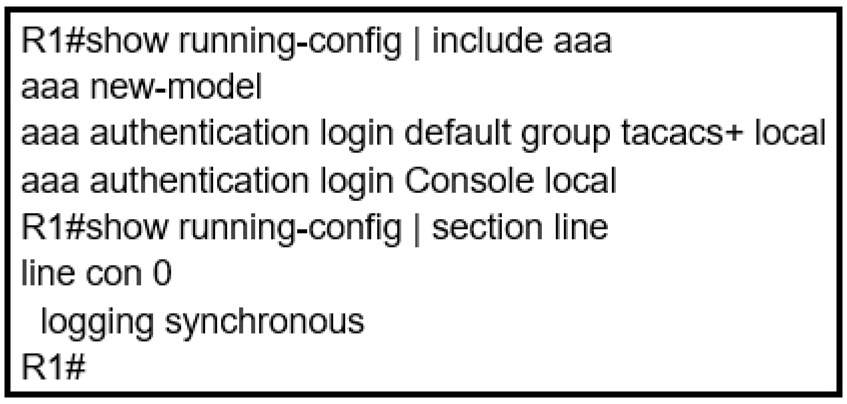

```
R1#show running-config | include aaa
aaa new-model
aaa authentication login default group tacacs+ local
aaa authentication login Console local
R1#show running-config | section line
line con 0
 logging synchronous
R1#
```

A. Add the aaa authentication login default none command to the global configuration.

B. Replace the capital ג€Cג€ with a lowercase ג€cג€ in the aaa authentication login Console local command.

C. Add the aaa authentication login default group tacacs+ local-case command to the global configuration.

D. Add the login authentication Console command to the line configuration

Correct Answer: *D*

Community vote distribution

D (100%)

⊟ **CraigB83** Highly Voted 3 years, 2 months ago

D sounds right:

Example 2: Console Access Using Line Password
Let's expand the configuration from Example 1 so that console login is only authenticated by the password set on line con 0.

The list CONSOLE is defined and then applied to line con 0.

We configure:

Router(config)# aaa authentication login CONSOLE line
In the command above:

the named list is CONSOLE.

there is only one authentication method (line).

Once a named list (in this example, CONSOLE) is created, it must be applied to a line or interface for it to come into effect. This is done using the login authentication list_name command:

Router(config)# line con 0
Router(config-line)# exec-timeout 0 0
Router(config-line)# password cisco
Router(config-line)# login authentication CONSOLE

https://www.cisco.com/c/en/us/support/docs/security-vpn/terminal-access-controller-access-control-system-tacacs-/10384-security.html

upvoted 9 times

⊟ **steiger** Most Recent 2 years ago

Selected Answer: D

D is right, the default authentication group is in use and you want it to use the Console group

upvoted 4 times

⊟ **error_909** 2 years, 3 months ago

The given answer is correct D

upvoted 1 times

⊟ **examShark** 2 years, 4 months ago

The given answer is correct

upvoted 1 times

⊟ **Wesgo** 2 years, 8 months ago

D is right. There are 2 authentication profiles here: (1) default and (2) Console. (1) will first authenticate with TACACS+ (that is mentioned by the question) and (2) has not been applied to the console. D is binding (2) to con 0 configuration.

upvoted 3 times

Question #129 *Topic 1*

Refer to the exhibit. An engineer is trying to connect to a device with SSH but cannot connect. The engineer connects by using the console and finds the displayed output when troubleshooting.

Which command must be used in configuration mode to enable SSH on the device?

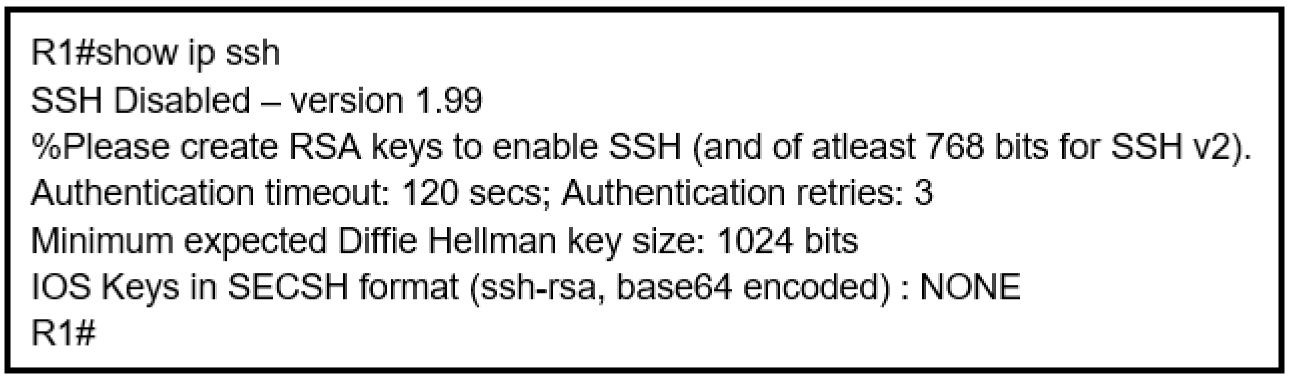

```
R1#show ip ssh
SSH Disabled – version 1.99
%Please create RSA keys to enable SSH (and of atleast 768 bits for SSH v2).
Authentication timeout: 120 secs; Authentication retries: 3
Minimum expected Diffie Hellman key size: 1024 bits
IOS Keys in SECSH format (ssh-rsa, base64 encoded) : NONE
R1#
```

A. no ip ssh disable

B. ip ssh enable

C. ip ssh version 2

D. crypto key generate rsa

Correct Answer: *D*

Community vote distribution

D (86%) 14%

Koume 11 months ago

Selected Answer: D

The log states that ssh is disabled because lacks of rsa key pairs. So this is the only answer here.

upvoted 2 times

mrnipsnips 1 year, 1 month ago

Selected Answer: D

D, you have to generate the key before enabling ssh

upvoted 3 times

quyle 1 year, 2 months ago

I test lab on eve -> D. Must crypto key generate rsa, then ip ssh version 2

upvoted 1 times

James1984 1 year, 5 months ago

Selected Answer: D

D is correct

upvoted 1 times

Nhan 1 year, 5 months ago

The message indicate that the rsa key is not yet generated. When you generate the rsa key using the command the ssh v1.99 will be enabled

upvoted 1 times

Mystic13 1 year, 7 months ago

You need to generate the rsa keys before you enable sshv2. D is correct

upvoted 1 times

Kimaf 1 year, 8 months ago

Selected Answer: C

Please read this from ENARSI book where it clearly says SSH enabled so how come our answer is D. It should be C.
SW1# show ip ssh
SSH Enabled - version 1.99 Authentication timeout: 120 secs; Authentication retries: 3 Minimum expected Diffie Hellman key size : 1024 bits IOS Keys in SECSH format(ssh-rsa, base64 encoded): ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAAAgQDtRqwdcEI+aGEXYmklh4G6pSJW1th6/Ivg4BCp19tO
BmdoW6NZahL2SxdzjKW8VIBjO1lVeaMfdmvKlpLjUlx7JDAkPs4Q39kzdPHY74MzD1/u+Fwvir8O5AQO
rUMkc5vuVEHFVc4WxQsxH4Q4Df10a6Q3UAOtnL4E0a7ez/imHw==

upvoted 1 times

Koume 11 months ago

Well i will cite the ENARSI book

"To check the version of SSH that is run□ning, use the show ip ssh command, as shown in Example 23-5. If it states version 1.99, it means versions 1 and 2 are running. If it states version 1, then SSHv1 is running, and if it states version 2, then SSHv2 is running."

upvoted 1 times

lcy1 1 year, 10 months ago

v1.99 means ssh v2 is enabled in config, but key is missing

upvoted 1 times

NH01 2 years, 1 month ago

The given answer is correct

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

Question #130 *Topic 1*

Which statement about IPv6 ND inspection is true?

A. It learns and secures bindings for stateless autoconfiguration addresses in Layer 3 neighbor tables.

B. It learns and secures bindings for stateless autoconfiguration addresses in Layer 2 neighbor tables.

C. It learns and secures bindings for stateful autoconfiguration addresses in Layer 3 neighbor tables.

D. It learns and secures bindings for stateful autoconfiguration addresses in Layer 2 neighbor tables.

Correct Answer: *B*

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/15-s/ip6f-15-s-book/ip6-snooping.pdf

Community vote distribution

B (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

Option B:
https://www.exam-answer.com/ipv6-nd-inspection-cisco-300-410-enarsi

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

Option B is correct:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/15-sy/ip6-nd-inspect.html

IPv6 ND Inspection
IPv6 ND inspection learns and secures bindings for stateless autoconfiguration addresses in Layer 2 neighbor tables. IPv6 ND inspection analyzes neighbor discovery messages in order to build a trusted binding table database, and IPv6 neighbor discovery messages that do not have valid bindings are dropped. A neighbor discovery message is considered trustworthy if its IPv6-to-MAC mapping is verifiable.

This feature mitigates some of the inherent vulnerabilities for the neighbor discovery mechanism, such as attacks on duplicate address detection (DAD), address resolution, device discovery, and the neighbor cache.

upvoted 1 times

Hurk2 11 months, 1 week ago

Selected Answer: B

B is correct

https://www.cisco.com/en/US/docs/ios-xml/ios/15-0se/features/ip6-snooping.html#GUID-5B40C0D5-3F0D-49FE-AA97-297F1B174BA9

upvoted 2 times

wts 1 year, 3 months ago

ND 2001:DB8:0:12::2 0017.5AED.7AF0 Gi0/2 1 0005 15s REACHABLE 288 s
- is this a Layer2 or Layer3 entry?
They will be independent of DHCP or SLAAC.

upvoted 1 times

Networkingguy 1 year, 11 months ago

IPv6 ND inspection learns and secures bindings for stateless autoconfiguration addresses in Layer 2 neighbor tables. IPv6 ND inspection analyzes neighbor discovery messages in order to build a trusted binding table database, and IPv6 neighbor discovery messages that do not have valid bindings are dropped. A neighbor discovery message is considered trustworthy if its IPv6-to-MAC mapping is verifiable.

This feature mitigates some of the inherent vulnerabilities for the neighbor discovery mechanism, such as attacks on duplicate address detection (DAD), address resolution, device discovery, and the neighbor cache.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/15-s/ip6f-15-s-book/ip6-snooping.pdf

upvoted 2 times

Networkingguy 1 year, 9 months ago

IPv6 ND inspection operates at Layer 2, or between Layer 2 and Layer 3, to provide IPv6 functions with security and scalability. Your software release may not support all the features documented in this module.

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #131 *Topic 1*

While troubleshooting connectivity issues to a router, these details are noticed:
✑ Standard pings to all router interfaces, including loopbacks, are successful.
✑ Data traffic is unaffected.
✑ SNMP connectivity is intermittent.
✑ SSH is either slow or disconnects frequently.
Which command must be configured first to troubleshoot this issue?

A. show policy-map control-plane

B. show policy-map

C. show interface | inc drop

D. show ip route

Correct Answer: *A*

GreatDane 1 year, 4 months ago
In this situation, data plane traffic (user traffic) is unaffected, while management plane traffic, such as SNMP and SSH, has problems. Troubleshooting must be made on control plane policing (CoPP).

A. show policy-map control-plane

Ref: Catalyst 6500 Release 15.0SY Software Configuration Guide

"Control Plane Policing (CoPP)
...
Monitoring CoPP

You can enter the show policy-map control-plane command for developing site-specific policies,
monitoring statistics for the control plane policy, and troubleshooting CoPP.
..."

Correct answer.

B. show policy-map

Wrong answer.

C. show interface | inc drop

Wrong answer.

D. show ip route

Wrong answer.
upvoted 3 times

Luvshah 2 months ago
Hi Grate Dane, can you please share your email address? Thanks
upvoted 1 times

error_909 2 years, 3 months ago
The given answer is correct
upvoted 2 times

examShark 2 years, 4 months ago
The given answer is correct
upvoted 2 times

Question #132

Topic 1

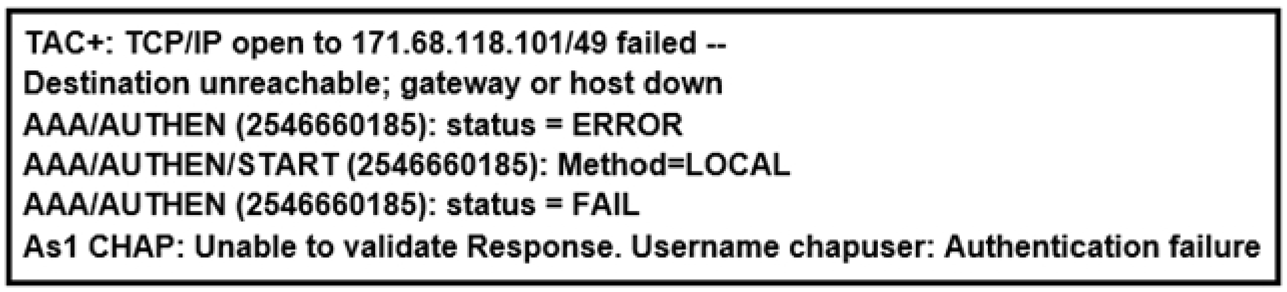

```
TAC+: TCP/IP open to 171.68.118.101/49 failed --
Destination unreachable; gateway or host down
AAA/AUTHEN (2546660185): status = ERROR
AAA/AUTHEN/START (2546660185): Method=LOCAL
AAA/AUTHEN (2546660185): status = FAIL
As1 CHAP: Unable to validate Response. Username chapuser: Authentication failure
```

Refer to the exhibit. Why is user authentication being rejected?

A. The TACACS+ server expects ג€userג€, but the NT client sends ג€domain/userג€.

B. The TACACS+ server refuses the user because the user is set up for CHAP.

C. The TACACS+ server is down, and the user is in the local database.

D. The TACACS+ server is down, and the user is not in the local database.

Correct Answer: *D*

Reference:

https://www.cisco.com/c/en/us/support/docs/security-vpn/terminal-access-controller-access-control-system-tacacs-/13864-tacacs-pppdebug.html

Community vote distribution

D (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

D is correct

upvoted 1 times

Noproblem22 1 year, 1 month ago

D is correct

upvoted 1 times

GreatDane 1 year, 4 months ago

The exhibit says it all:

TAC+...gateway or host down
AAA/...Method=LOCAL
AAA/...status=FAIL

A. The TACACS+ server expects “user”, but the NT client sends “domain/user”.

Wrong answer.

B. The TACACS+ server refuses the user because the user is set up for CHAP.

Wrong answer.

C. The TACACS+ server is down, and the user is in the local database.

Wrong answer.

D. The TACACS+ server is down, and the user is not in the local database.

Correct answer.

upvoted 2 times

AliMo123 2 years, 1 month ago

server is down bc 171.68.118.101/49 failed
also local (2546660185) status is fail
D is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

Question #133 *Topic 1*

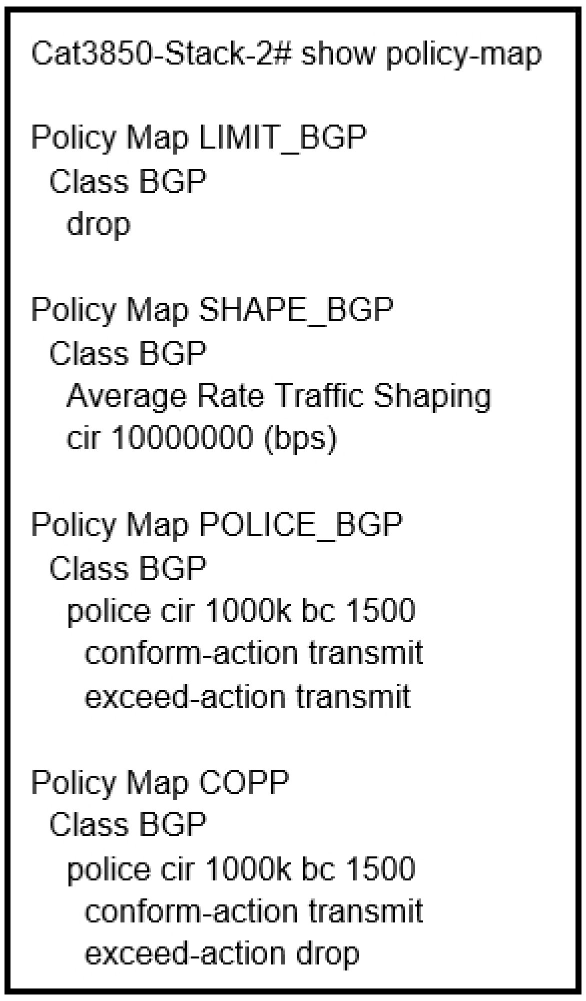

```
Cat3850-Stack-2# show policy-map

Policy Map LIMIT_BGP
  Class BGP
    drop

Policy Map SHAPE_BGP
  Class BGP
    Average Rate Traffic Shaping
    cir 10000000 (bps)

Policy Map POLICE_BGP
  Class BGP
    police cir 1000k bc 1500
      conform-action transmit
      exceed-action transmit

Policy Map COPP
  Class BGP
    police cir 1000k bc 1500
      conform-action transmit
      exceed-action drop
```

Refer to the exhibit. Which control plane policy limits BGP traffic that is destined to the CPU to 1 Mbps and ignores BGP traffic that is sent at higher rate?

A. policy-map SHAPE_BGP

B. policy-map LIMIT_BGP

C. policy-map POLICE_BGP

D. policy-map COPP

Correct Answer: *D*

Community vote distribution

D (60%) C (40%)

guy276465281819372 Highly Voted 5 months, 4 weeks ago

it would have been nice if cisco would use professional terminology and not use a word like "ignore". annoying!

upvoted 5 times

JieW Most Recent 4 months, 1 week ago

Selected Answer: D

Looking at the actual configs, POLICE_BGP is just a name and not actually doing "policing".
COPP is though.
Vote D

upvoted 1 times

guy276465281819372 4 months, 2 weeks ago

Selected Answer: C

i think ignore means let it pass through

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

D SI CORRECT

upvoted 1 times

Hurk2 11 months, 1 week ago

Selected Answer: D

Di is correct, C does not police the exceeding traffic

upvoted 1 times

smithkeith0023366 1 year ago

Selected Answer: D

Voting. policy-map COPP is correct.

upvoted 1 times

Noproblem22 1 year, 1 month ago

I believe "ignore" the traffic that exceeds means "drop", in this case D is coorect

upvoted 2 times

CisconAWSGURU 1 year, 1 month ago

Selected Answer: C

Correct is C

upvoted 1 times

Remsync 1 year, 2 months ago

Selected Answer: D

Correct is D

upvoted 2 times

Samurai55_1998_01 1 year, 2 months ago

I believe, in this context, "ignore" means discard so the answer is "D".

upvoted 1 times

NoUserName1234 1 year, 3 months ago

quick check :
https://networklessons.com/quality-of-service/policing-configuration-example
Exceed action transmit is 'if nothing else is claiming bw then you may proceed'
Exceed action drop is 'you're out of luck if you go above cir' = dropped packets
Answer D- COPP is correct

upvoted 1 times

wts 1 year, 3 months ago

Selected Answer: C

police cir 1000k - "policy limits BGP traffic that is destined to the CPU to 1 Mbps"
exceed-action transmit - "ignores BGP traffic that is sent at higher rate"

upvoted 2 times

> **Remsync** 1 year, 2 months ago
>
> C is wrong.
> Should be "exceed-action drop". Any traffic that goes beyond 1Mbps should be dropped, not transmited.
>
> upvoted 2 times

GreatDane 1 year, 4 months ago

“limits BGP traffic that is destined to the CPU”
class BGP

“to 1 Mbps”
police cir 1000k...
conform-action transmit

“ignores BGP traffic that is sent at higher rate”
exceed-action drop

It’s policy-map COPP.

upvoted 3 times

Hack4 1 year, 8 months ago

The given answer is correct

upvoted 1 times

wts 1 year, 10 months ago

Limits and ignores at the same time? Can anyone explain this?

upvoted 1 times

> **wts** 1 year, 10 months ago
>
> The authors do not want to say that to ignore is to discard?
>
> upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #134 *Topic 1*

Which statement about IPv6 RA Guard is true?

A. It does not offer protection in environments where IPv6 traffic is tunneled.

B. It cannot be configured on a switch port interface in the ingress direction.

C. Packets that are dropped by IPv6 RA Guard cannot be spanned.

D. It is not supported in hardware when TCAM is programmed.

Correct Answer: *A*

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-16/ip6f-xe-16-book/ip6-ra-guard.pdf

Community vote distribution

A (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: A

Correct: A :

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-3s/ip6f-xe-3s-book/ip6-ra-guard.pdf

upvoted 1 times

GreatDane 1 year, 4 months ago

Ref: IPv6 First-Hop Security Configuration Guide, Cisco IOS XE Release 3S

“CHAPTER 1
IPv6 RA Guard
...
Restrictions for IPv6 RA Guard
...
• The IPv6 RA Guard feature does not offer protection in environments where IPv6 traffic is tunneled.
...”

A. It does not offer protection in environments where IPv6 traffic is tunneled.

Correct answer.

B. It cannot be configured on a switch port interface in the ingress direction.

Wrong answer.

C. Packets that are dropped by IPv6 RA Guard cannot be spanned.

Wrong answer.

D. It is not supported in hardware when TCAM is programmed.

Wrong answer.

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

murinha10 2 years, 5 months ago

The correct answer is A.

upvoted 1 times

dk1996 2 years, 7 months ago

C is the correct !!!

upvoted 1 times

Pb1805 2 years, 7 months ago

C is definately wrong. Packets dropped by the IPv6 RA Guard feature can be spanned
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/15-s/ip6f-15-s-book/ip6-ra-guard.pdf

upvoted 4 times

Question #135 Topic 1

An engineer must configure a Cisco router to initiate secure connections from the router to other devices in the network but kept failing.
Which two actions resolve the issue? (Choose two.)

A. Configure transport input ssh command on the console.

B. Configure a domain name.

C. Configure a crypto key to be generated.

D. Configure a source port for the SSH connection to initiate.

E. Configure a TACACS+ server and enable it.

Correct Answer: *BC*

Community vote distribution

BC (75%) AD (25%)

AlexInShort12 3 days, 12 hours ago

Selected Answer: AD

I'm with Pietjeplukgeluk.
The question seems to be indicating that the router is not able to make a SSH connection OUT. A-D could be the reason.

upvoted 1 times

Pietjeplukgeluk 3 weeks, 5 days ago

This question is utterly stupid as it seems to indicate the Cisco router only acts as a SSH client. An SSH client does not require ANY configuration.
Again, B+C seems correct if you look at the options, but the question itself seems a bit strange.

upvoted 2 times

GreatDane 1 year, 4 months ago

Ref: Configuring Secure Shell on Routers and Switches Running Cisco IOS – Cisco

“…
Set Up an IOS Router or Switch as SSH Client

There are four steps required to enable SSH support on a Cisco IOS router:

1. Configure the hostname command.

2. Configure the DNS domain.

3. Generate the SSH key to be used.

4. Enable SSH transport support for the virtual type terminal (vtys).
…”

A. Configure transport input ssh command on the console.

Wrong answer.

B. Configure a domain name.

Correct answer.

C. Configure a crypto key to be generated.

Correct answer.

D. Configure a source port for the SSH connection to initiate.

Wrong answer.

E. Configure a TACACS+ server and enable it.

Wrong answer.

upvoted 2 times

Carl1999 1 year, 10 months ago

Selected Answer: BC

B,C
need these commands to configure ssh.

upvoted 2 times

Carl1999 1 year, 10 months ago

#hostname
#ip domain-name
#crypto key generate rsa

upvoted 1 times

JingleJangus 1 year, 10 months ago

Selected Answer: BC

BC are correct.

A. makes no sense since the console isnt able to be accessed via ssh or telnet.

D. isnt it because source ports are autogenerated and dont need to be explicitly configured.

E. Unless the user needs authorization, this answer makes no sense.

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #136 *Topic 1*

When configuring Control Plane Policing on a router to protect it from malicious traffic, an engineer observes that the configured routing protocols start flapping on that device.
Which action in the Control Plane Policy prevents this problem in a production environment while achieving the security objective?

A. Set the conform-action and exceed-action to transmit initially to test the ACLs and transmit rates and apply the Control Plane Policy in the output direction.

B. Set the conform-action and exceed-action to transmit initially to test the ACLs and transmit rates and apply the Control Plane Policy in the input direction.

C. Set the conform-action to transmit and exceed-action to drop to test the ACLs and transmit rates and apply the Control Plane Policy in the input direction.

D. Set the conform-action to transmit and exceed-action to drop to test the ACLs and transmit rates and apply the Control Plane Policy in the output direction.

Correct Answer: *B*

Community vote distribution

B (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

B correct:

https://www.exam-answer.com/configure-control-plane-policing-prevent-routing-protocol-flapping

upvoted 1 times

HungarianDish 7 months, 4 weeks ago

Selected Answer: B

I agree with the post from Networkingguy, first we permit (transmit) all traffic to see how much packets are exceeding. Pls see:
https://networklessons.com/cisco/ccie-routing-switching-written/copp-control-plane-policing

However, we would need to use exceed-action drop in order to protect the control plane (security objective). The question is formed ambiguously. Still I vote for B, because testing should be performed before setting the drop action.

upvoted 3 times

chris7890 1 year, 2 months ago

can someone resolve whether answer B or C are correct? Thanks

upvoted 1 times

JOKERR 1 year, 6 months ago

I think given answer is right. This is an excerpt from Cisco:

he CoPP feature on a Cisco device does exactly what it sounds like: It polices the traffic coming to the control plane. For this purpose, the control plane is treated as a logical source and destination, with its own inbound and outbound interfaces. Only traffic that is destined for the control plane is policed as part of this feature. This is in addition to any policing, filtering, or any other processing done at the interface where the packet was received by the device.

So, you police traffic coming to the Control Plane so that it doesn't have to process it.

https://www.ciscopress.com/articles/article.asp?p=2928193&seqNum=3

upvoted 1 times

Kimaf 1 year, 8 months ago

I know the answer is either A or B because of the ACL but here is the a paragraph from the OCG Enarsi book page 861
Direction: CoPP can be applied to packets entering or leaving the control plane interface. Therefore, the correct direction needs to be specified. For incoming packets, you specify input, and for outgoing packets you specify output. Direction can be verified with the output of show policy-map control-plane as well. Note that not all versions support output CoPP, and for the ones that do, you need to ensure that the correct traffic is being classified in the ACLs and the class maps. For example, when it comes to BGP, OSPF (Open Shortest Path First), and EIGRP, you typically use output CoPP for the replies that are being sent because of an already received packet. For ICMP, it would be error and informational reply messages. For Telnet, SSH (Secure Shell), HTTP (Hypertext Transfer Protocol), or SNMP (Simple Network Management Protocol), you would be dealing with replies or traps. If the ACL and class map are not configured appropriately for the replies, the desired result will not be achieved.
So my guess is A.

upvoted 1 times

[Removed] 4 months ago

I also viewed this excerpt as the answer, but the question is talking about protecting the router from malicious traffic, and this (to me) meant inbound traffic is being policed and maybe some of the routing protocol packets are getting caught in the policy map

upvoted 1 times

Carl1999 1 year, 10 months ago

B or C correct.
I only know that" the input direction" is correct.

upvoted 1 times

Networkingguy 1 year, 9 months ago

Input direction because we are sussing out Malicious public traffic that might come in, and we are testing so we would want to use conform and exceed to just give results of what we are working with.

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Networkingguy 1 year, 11 months ago

ExamShark, you are a twat for copy and pasting the same response on every question. I haven't seen you say anything useful, hope you get the lot ya dawg

upvoted 12 times

Question #137 *Topic 1*

In which two ways does the IPv6 First-Hop Security Binding Table operate? (Choose two.)

A. by IPv6 HSRP to make sure neighbors are authenticated before being used as gateways

B. by various IPv6 guard features to validate the data link layer address

C. by the recovery mechanism to recover the binding table in the event of a device reboot

D. by IPv6 routing protocols to securely build neighborships without the need of authentication

E. by storing hashed keys for IPsec tunnels for the built-in IPsec features

Correct Answer: *BC*

Community vote distribution

BC (100%)

studybuddy10 Highly Voted 2 years, 1 month ago

given answer is correct, first two lines from this article:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/15-s/ip6-fhs-bind-table.html

upvoted 7 times

inteldarvid Most Recent 5 months, 2 weeks ago

Selected Answer: BC

B, C :

https://www.exam-answer.com/ipv6-first-hop-security-binding-table-operation

upvoted 1 times

GreatDane 1 year, 4 months ago

Ref: IPv6 First-Hop Security Binding Table – Cisco

"...
Overview of the IPv6 First-Hop Security Binding Table
...
This database, or binding table, is used by various IPv6 guard features to validate the link-layer address (LLA), the IPv4 or IPv6 address, and the prefix binding of the neighbors to prevent spoofing and redirect attacks.
...
IPv6 First-Hop Security Binding Table Recovery Mechanism

The IPv6 first-hop security binding table recovery mechanism enables the binding table to recover in the event of a device reboot.
..."

A. by IPv6 HSRP to make sure neighbors are authenticated before being used as gateways

Wrong answer.

B. by various IPv6 guard features to validate the data link layer address

Correct answer.

C. by the recovery mechanism to recover the binding table in the event of a device reboot

Correct answer.

D. by IPv6 routing protocols to securely build neighborships without the need of authentication

Wrong answer.

E. by storing hashed keys for IPsec tunnels for the built-in IPsec features

Wrong answer.

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 3 times

Question #138 *Topic 1*

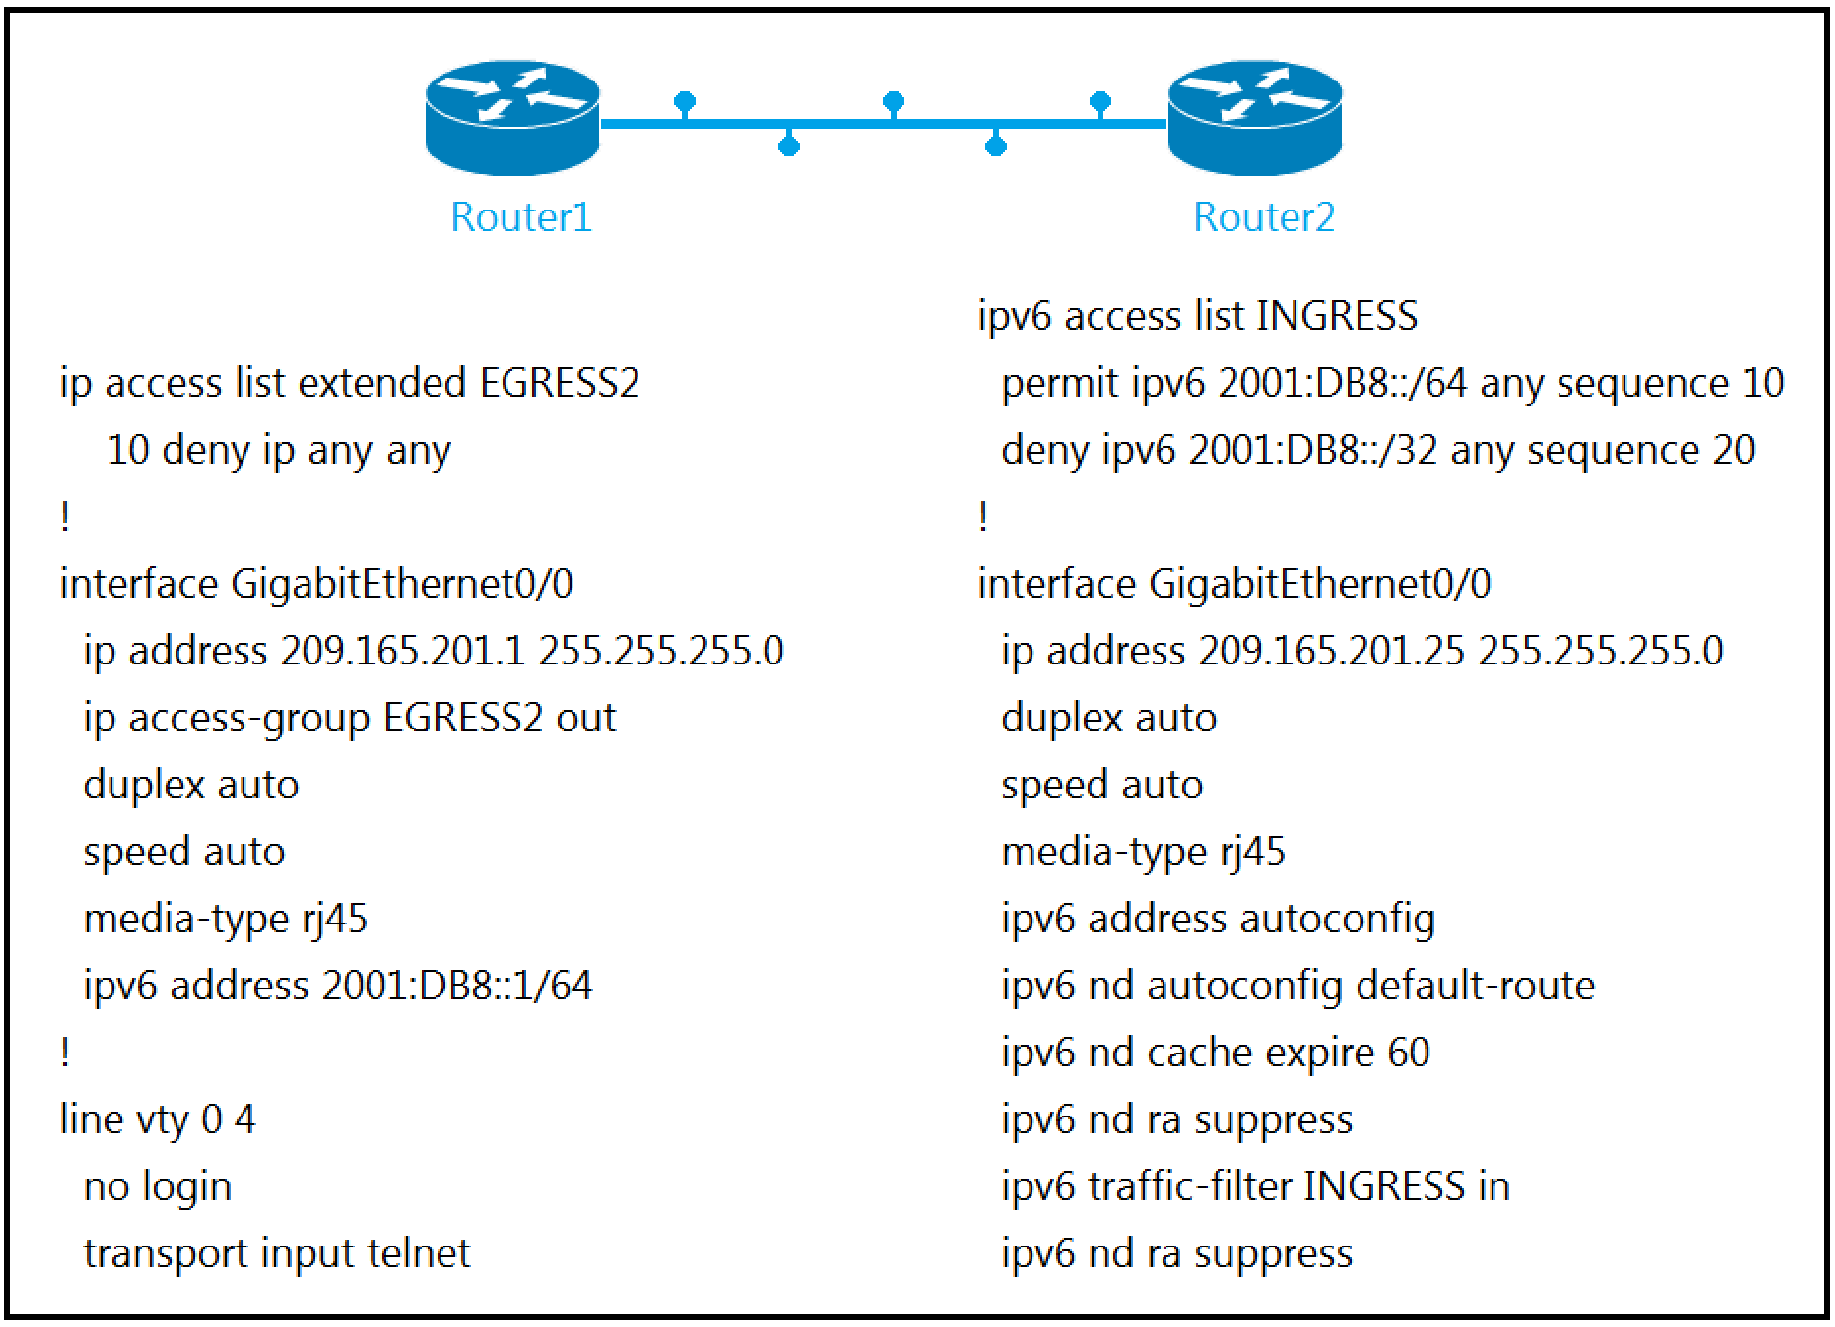

```
ip access list extended EGRESS2
   10 deny ip any any
!
interface GigabitEthernet0/0
  ip address 209.165.201.1 255.255.255.0
  ip access-group EGRESS2 out
  duplex auto
  speed auto
  media-type rj45
  ipv6 address 2001:DB8::1/64
!
line vty 0 4
  no login
  transport input telnet
```

```
ipv6 access list INGRESS
  permit ipv6 2001:DB8::/64 any sequence 10
  deny ipv6 2001:DB8::/32 any sequence 20
!
interface GigabitEthernet0/0
  ip address 209.165.201.25 255.255.255.0
  duplex auto
  speed auto
  media-type rj45
  ipv6 address autoconfig
  ipv6 nd autoconfig default-route
  ipv6 nd cache expire 60
  ipv6 nd ra suppress
  ipv6 traffic-filter INGRESS in
  ipv6 nd ra suppress
```

Refer to the exhibit. The engineer configured and connected Router2 to Router1. The link came up but could not establish a Telnet connection to Router1 IPv6 address of 2001:DB8::1.

Which configuration allows Router2 to establish a Telnet connection to Router1?

A. ipv6 unicast-routing

B. permit ICMPv6 on access list INGRESS for Router2 to obtain IPv6 address

C. permit ip any any on access list EGRESS2 on Router1

D. IPv6 address on GigabitEthernet0/0

Correct Answer: *C*

Community vote distribution

D (45%) B (32%) C (19%) 3%

MP_iBGP Highly Voted 2 years, 2 months ago

Correct answer is B because when R1 will send nd ra to R2 for its autoconfig, its access-list INGRESS will drop it.
LAB for test :

```
R2#show ipv6 access-list
IPv6 access list INGRESS
permit ipv6 2001:DB8::/64 any (1 match) sequence 10
deny ipv6 2001:DB8::/32 any sequence 20
permit icmp any any (5 matches) sequence 30

R2#telnet 2001:db8::1
Trying 2001:DB8::1 ... Open

R1>
```

upvoted 16 times

donjime 2 years, 2 months ago

RA are suppressed by the comand ipv6 nd ra suppress on the interface

upvoted 1 times

[Removed] 1 year, 11 months ago

You're right.. It stops that router from advertising but it doesnt stop it from responding to RA messages.. Add the icmp to the acl and it will be able to generate an ipv6 address since autoconfig is enabled. I also labbed to verify...

upvoted 7 times

asans 1 year, 8 months ago

B is correct, permitting icmp on R2 enables it to receive RA with the prefix info and thus generate an IPv6 address. D works but the key here is to use the ipv6 address autoconfig feature rather the manual IPv6 address

upvoted 3 times

wts 1 year, 8 months ago

What message exactly contains address 2001:DB8::/32 in the source and what does it matter if what is forbidden is allowed by the line above?

All of these messages should use link-local addresses (FE80::/64) as their source. I believe the results of your test, but how to explain it?

upvoted 1 times

lcy1 Highly Voted 👍 1 year, 10 months ago

tested in lab - A doesn't work, unless B is done. B by itself doesn't help without A
D helps instantly.
So it depends how many answers cisco wants on real exam - if one, then it is D, if two, then it is AB

upvoted 8 times

samael666 Most Recent ⏱ 1 month, 2 weeks ago

Correct answer is D.
A. it says the link came up, so is enable by default
B. on IPv6 ACLs is enabled by default
C. it has nothing to do with it
D. is the only choice, but consider that there is a autonconfig command so withouht this it will work as well.

upvoted 1 times

guy276465281819372 4 months, 2 weeks ago

Selected Answer: D

D would solve this question in instant

upvoted 1 times

sgtmajvimy 4 months, 3 weeks ago

Selected Answer: B

B is correct, its configured for autoconfig, the ACL blocks R2 from getting the RA from R1.

upvoted 1 times

inteldarvid 5 months ago

Selected Answer: D

sorry my answer before, I thinking about this question for a while, and the correct answer is "D" and not "B". The key command is "ipv6 nd ra suppress" we are blocking RA ads on IPV6 and an ACL that allows ICMPv6 is not needed we are already blocking it. It's option "D"

upvoted 1 times

MicMillon 5 months, 2 weeks ago

Selected Answer: B

B is correct

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

option B is correct:

https://docs.ruckuswireless.com/fastiron/08.0.60/fastiron-08060-securityguide/GUID-4F7DBEAC-7D2F-4FE2-86A8-94C376D63B2E.html

upvoted 1 times

MicMillon 5 months, 2 weeks ago

Selected Answer: B

correct answer is B. its not C because thats only blocking ipv4, and its not D because its using auto-discovery to assign v6 address

upvoted 1 times

Malasxd 7 months, 1 week ago

Selected Answer: B

I would chose "B".

Nothin works without "A", but we don't know whether it was inserted or not in both routers.

C is definily not right. EGRESS2 is a IPv4 ACL and it's does not works for IPv6 packets.

D Would not work because R2 would need use NDP to discover R1's MAC address, and NDP works with ICMP that is blocked by INGRESS ACL.

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

"D" actually works. Test it. Setting an ipv6 address manually is enough for telnet to work. permit ICMPv6 is not necessary in this case, as NDP is not used for ipv6 address configuration here.
upvoted 1 times

Malasxd 7 months, 1 week ago
I forgot to mention one thing.

The address of NDP and RS/RA packets are link-local address. Because of that the INGRESS ACL does not allow them in sequence 10.
upvoted 1 times

HungarianDish 6 months, 3 weeks ago
A) #ipv6 unicast-routing -> Yes, I agree, normally it should be enabled first. Stil, setting ipv6 addresses manually is enough for a basic communication between directly connected neighbors. Just test it.
B) permit ICMPv6 -> It is not needed if the ipv6 address is already configured manually. Setting an ipv6 address is enough for telnet to work.
upvoted 2 times

HungarianDish 7 months, 3 weeks ago
Selected Answer: D
Answer A + B or Answer D.
A) We need to configure #ipv6 unicast-routing on R1, so it can start to send RA messages on the local segment.
+
B) permit ICMPv6 on access list INGRESS on R2
-> My assumption was that ipv6 acl implicit rules contain permition for ICMPv6 neighbor discovery protocol.
I also read it on cisco learning network that these implicit entries exist at the end of each IPv6 ACL to allow neighbour discovery.
Then I labbed this scenario in CML, and it turned out that in this case I need to explicitly add these lines to the ACL for NDP to work well.
(At least on that IOS in CML.)

```
permit icmp any any nd-na
permit icmp any any nd-ns
permit icmp any any router-advertisement
permit icmp any any router-solicitation
```

D) IPv6 address on GigabitEthernet0/0 -> The workaround if only one answer can be chosen.
upvoted 6 times

Hurk2 11 months, 1 week ago
Selected Answer: A
I have labed this, telnet works from R2 to R1 with exactly the same configuration when I enable ipv6 unicast-routing. So A is correct
upvoted 1 times

DUBC89x 1 year ago
Selected Answer: B
MP_iBGP is correct. I also used a LAB and verified results.
Debug

```
*Dec 5 23:42:07.442: [IPv6 Input]IPv6RT[default]: ND, Added path FE80::C804:CFF:FEFE:1C/GigabitEthernet1/0 (A:0x1/F:0x0)
*Dec 5 23:42:07.446: [IPv6 Input]IPv6RT[default]: ND, Route add 2001:DB8::/64 [new 2/0]
*Dec 5 23:42:07.450: [IPv6 Input]IPv6RT[default]: ND, Added path ::/GigabitEthernet1/0 (A:0x1/F:0x0)
*Dec 5 23:42:07.458: [IPv6 RIB Event Handler]IPv6RT[default]: Event: ::/0, Add, owner ND, previous None
*Dec 5 23:42:07.466: [IPv6 RIB Event Handler]IPv6RT[default]: Event: 2001:DB8::/64, Add, owner ND, previous None
```

upvoted 2 times

Edwinmolinab 1 year, 1 month ago
Selected Answer: B
To obtain an IPv6 address a client must be enable to receive icmpv6 particularly RA and to avoid duplicate address NA. B is the correct answer. ICMPv6 to avoid local troubles default enable values must be permitted.
upvoted 1 times

TECH3K3 1 year, 5 months ago
Selected Answer: B
Answer B
I wasn't sure, so I lab it myself as so many conflicting replies.
I permitted icmp to the INGRESS acl and was R2 got an IPv6 address, and I was able to Telnet to R1.
upvoted 2 times

Nhan 1 year, 5 months ago
It can't not be d because the interface g0/0 has an ipv6 address already the command on the interface is ipv6 address autoconfig. The answer is C
upvoted 1 times

WAKIDI 1 year, 5 months ago
Selected Answer: D
D is correct. I Labed it using Cisco Packet Tracer. D is proven. A,B and C is wrong. B is wrong even INGRESS ACL is not used.

```
R2(config-if)#ipv6 address 2001:db8::2/64
R2#telnet 2001:db8::1
Trying 2001:DB8::1 ...Open

R1>
```

upvoted 2 times

Question #139 *Topic 1*

An engineer configured Reverse Path Forwarding on an interface and noticed that the routes are dropped when a route lookup fails on that interface for a prefix that is available in the routing table.
Which interface configuration resolves the issue?

A. ip verify unicast source reachable-via l2-src

B. ip verify unicast source reachable-via allow-default

C. ip verify unicast source reachable-via any

D. ip verify unicast source reachable-via rx

Correct Answer: *C*

Community vote distribution

C (100%)

⊟ **inteldarvid** 5 months, 2 weeks ago

Selected Answer: C

correct is C :

https://www.exam-answer.com/configure-reverse-path-forwarding

upvoted 1 times

⊟ **GreatDane** 1 year, 4 months ago

Ref: Security - Configuring Network Security [Support] - Cisco Systems

"...
Configuring the Unicast RPF Check Mode

There are two Unicast RPF check modes:

• Strict check mode, which verifies that the source IP address exists in the FIB table and verifies that the source IP address is reachable through the input port.
• Exist-only check mode, which only verifies that the source IP address exists in the FIB table.
...
When configuring the Unicast RPF check mode, note the following information:

• Use the rx keyword to enable strict check mode.
• Use the any keyword to enable exist-only check mode.
• Use the allow-default keyword to allow use of the default route for RPF verification.
..."

The route lookup failed, but the prefix is in the routing table. RPF Exist-only check mode is the way to go.

A. ip verify unicast source reachable-via l2-src

Wrong answer.

B. ip verify unicast source reachable-via allow-default

Wrong answer.

C. ip verify unicast source reachable-via any

Correct answer.

D. ip verify unicast source reachable-via rx

Wrong answer.

upvoted 4 times

⊟ **Luvshah** 2 months ago

Hi, Could you please provide me your email address ? Thanks

upvoted 1 times

⊟ **GReddy2323** 9 months, 3 weeks ago

Thank you very much for your awesome answers.

upvoted 1 times

⊟ **wts** 1 year, 9 months ago

the packet is dropped even though there is a route for the source address in the routing table - seems so much clearer what's going on

upvoted 1 times

⊟ **Hack4** 1 year, 10 months ago

THE given answer is correct

upvoted 1 times

Question #140 Topic 1

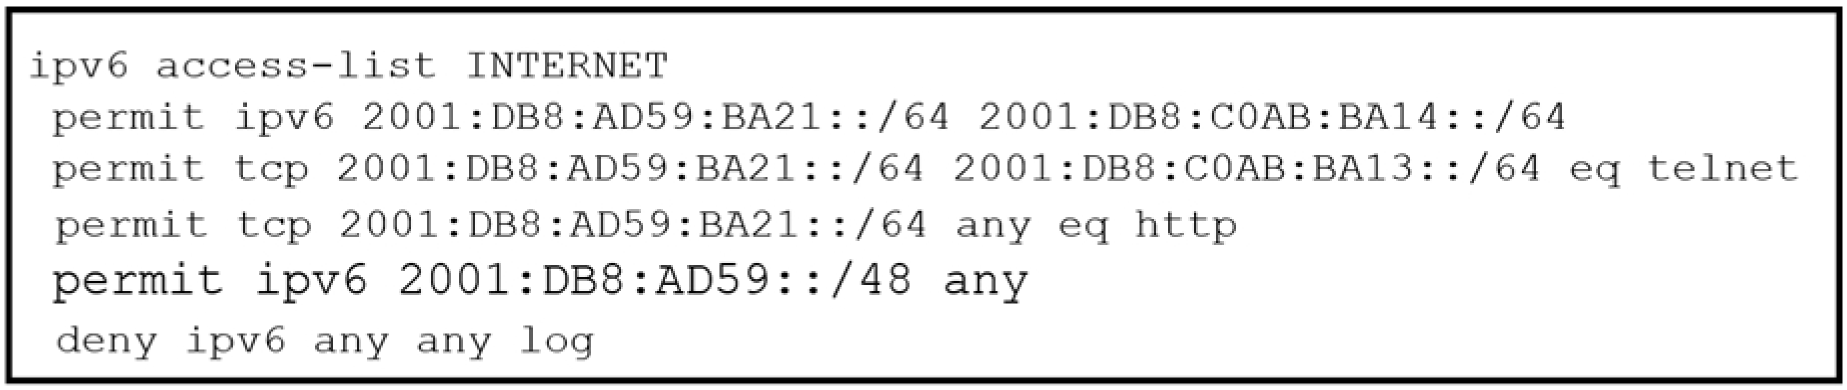

```
ipv6 access-list INTERNET
 permit ipv6 2001:DB8:AD59:BA21::/64 2001:DB8:C0AB:BA14::/64
 permit tcp 2001:DB8:AD59:BA21::/64 2001:DB8:C0AB:BA13::/64 eq telnet
 permit tcp 2001:DB8:AD59:BA21::/64 any eq http
 permit ipv6 2001:DB8:AD59::/48 any
 deny ipv6 any any log
```

Refer to the exhibit. When monitoring an IPv6 access list, an engineer notices that the ACL does not have any hits and is causing unnecessary traffic through the interface

Which command must be configured to resolve the issue?

A. ip access-group INTERNET in

B. ipv6 traffic-filter INTERNET in

C. ipv6 access-class INTERNET in

D. access-class INTERNET in

Correct Answer: *C*

Reference:

https://www.cisco.com/c/en/us/support/docs/ip/ip-version-6/113126-ipv6-acl-00.html

Community vote distribution

B (79%) C (21%)

Mishranihal737 2 months, 2 weeks ago

Selected Answer: B

It's asking for interface that's why traffic-filter. Access-class is used for control plane.

upvoted 1 times

Brand 3 months, 4 weeks ago

Selected Answer: B

R1(config-if)#ipv6 traffic-filter ?
WORD Access-list name

R1(config-if)#ipv6 traffic-filter

upvoted 1 times

sgtmajvimy 4 months, 3 weeks ago

Selected Answer: B

i concur, its B

upvoted 1 times

Chiaretta 5 months, 1 week ago

Selected Answer: B

Answer is B

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

the answer corret is B:
Line vty: acces-class
line interface: traffic-filter

https://www.cisco.com/c/en/us/td/docs/ios/ipv6/command/reference/ipv6_book/ipv6_05.html#wp2274594

upvoted 1 times

sajjad_gayyem 5 months, 3 weeks ago

Selected Answer: C

Im going with C, hence its denied and permitted the telnet traffics, so this ACl must be applied under the VTY lines, so for VTY line we must use
Applying an IPv6 ACL to the Virtual Terminal Line
SUMMARY STEPS
1. enable
2. configure terminal

3. line [aux| console| tty| vty] line-number[ending-line-number]
4. ipv6 access-class ipv6-access-list-name {in| out}
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_acl/configuration/xe-16/sec-data-acl-xe-16-book/ip6-acls-xe.html
upvoted 2 times

Koume 11 months ago

Selected Answer: B

B, because is talking about incesary traffic for the interface. Access Class is for apply Line vty ACL.
upvoted 1 times

Brand 4 months, 2 weeks ago

So line vty isn't an interface?
upvoted 1 times

Noproblem22 1 year, 1 month ago

B is right
https://community.cisco.com/t5/network-security/ipv6-access-class-vs-ipv6-traffic-filter/td-p/1510357#:~:text=The%20%27ipv6%20access-class%27%20command%20is%20used%20to%20filter,%28i.e.%20management%20traffic%29.%20Command%20reference%20%28with%20example%29%3A%20http%3A%2F%2Fwww.cisco.com%2Fen%2FUS%2Fdocs%2Fios%2Fipv6%2Fcommand%2Freference%2Fipv6_05.html%23wp2274594
upvoted 1 times

CisconAWSGURU 1 year, 1 month ago

Selected Answer: B

Answer is B
upvoted 1 times

mrnipsnips 1 year, 1 month ago

Selected Answer: C

Traffic filter
upvoted 1 times

Kapoduster 1 year, 2 months ago

Selected Answer: B

B is correct. :

```
R2(config-if)#ipv6 traff?
traffic-filter

R2(config-if)#ipv6 acces?
% Unrecognized command
R2(config-if)#ipv6 acces
```

upvoted 1 times

jarz 1 year, 2 months ago

Selected Answer: B

traffic-filter
upvoted 2 times

jarz 1 year, 2 months ago

Selected Answer: B

As AliMo123 says
upvoted 1 times

MasterP007 1 year, 2 months ago

C - is Incorrect. There's no access-class in IPv6

```
R4(config-if)#ipv6 access-class INTERNET in
^
```

upvoted 2 times

NoUserName1234 1 year, 3 months ago

When reading the mentioned link it's clear that it's answer B, as Alimo also states
upvoted 1 times

Duck2Duck 1 year, 5 months ago

Bad ipv6 acl... breaks neighbor discovery. Otherwise..B
upvoted 1 times

AliMo123 2 years, 1 month ago

B is correct
"In order to assign an IPv6 ACL to an interface, use this command in interface configuration mode: ipv6 traffic-filter access-list-name {in | out}"
upvoted 4 times

Question #141 Topic 1

Which configuration feature should be used to block rogue router advertisements instead of using the IPv6 Router Advertisement Guard feature?

A. VACL blocking broadcast frames from nonauthorized hosts

B. PVLANs with promiscuous ports associated to route advertisements and isolated ports for nodes

C. PVLANs with community ports associated to route advertisements and isolated ports for nodes

D. IPv4 ACL blocking route advertisements from nonauthorized hosts

Correct Answer: *B*

Community vote distribution

B (80%) D (20%)

Dirkd0344 Highly Voted 2 years ago

The answer is not D, as this is regarding IPv6. The answer would be B. You would configure the switch with PVLANs, configure the switchport where you would expect to see RAs as a promiscuous port, and configure the client ports as isolated ports. With this configuration if any rogue RAs came in on an isolated port it would not be able to offer SLAAC addresses to any other client on the other isolated ports.

upvoted 10 times

baid 1 year, 10 months ago

Thanks for your explanation. It's right.

upvoted 2 times

chris110 Most Recent 3 months, 3 weeks ago

Selected Answer: B

To block rogue router advertisements in an IPv6 network, you should use option B:

B. PVLANs (Private VLANs) with promiscuous ports associated with route advertisements and isolated ports for nodes.

Private VLANs help in segmenting traffic within a VLAN and provide isolation between devices within the same VLAN. In this context, you can configure a PVLAN such that the promiscuous port (connected to a trusted router) is allowed to send router advertisements, while the isolated ports (connected to end-user devices) are not allowed to send such advertisements. This way, you can prevent rogue router advertisements from unauthorized sources within the same VLAN.

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

B option:

https://www.exam-answer.com/which-configuration-feature-blocks-rogue-router-advertisements-ipv6

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: B

https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2020/pdf/BRKSEC-3200.pdf
Mitigating Rogue RA: Host Isolation
Private VLANs (PVLAN) where nodes (isolated port) can only contact the official router (promiscuous port)

upvoted 1 times

GreatDane 1 year, 4 months ago

Ref: Advanced IPv6 Security Threats and Mitigation – Cisco

“LAN Security with First Hop Security (FHS)
…
Mitigating Rogue RA: Host Isolation

Prevent Node-Node Layer-2 communication by using:

• Private VLANs (PVLAN) where nodes (isolated port) can only contact the official router (promiscuous port)
…”

A. VACL blocking broadcast frames from nonauthorized hosts

Wrong answer.

B. PVLANs with promiscuous ports associated to route advertisements and isolated ports for nodes

Correct answer.

C. PVLANs with community ports associated to route advertisements and isolated ports for nodes

Wrong answer.

D. IPv4 ACL blocking route advertisements from nonauthorized hosts

Wrong answer.
upvoted 2 times

PrettyStupid 1 year, 1 month ago
Agreed with GreatDane, checked the session video from cisco live (min 09:25 to 11:40 aprox) https://www.youtube.com/watch?v=RCxC2gIV4jo
upvoted 1 times

kellyDD 1 year, 6 months ago
promiscuous ports and isolated ports can communicate, right?
upvoted 1 times

thanh123 1 year, 8 months ago
Selected Answer: B
Techincally, you can use VACL to block RA but there are some issues. I haven't tested because GNS3 won't support VACL or private VLAN, I even don't have physical hardware, either. So correct me if I'm wrong:
1. You can use ACL to filter IP or MAC of rouge host generates RA. Downside of this is that if rouge router change IP or MAC, you have to change the ACL as well, which is not scale very well
2. If we choose to filter based on Layer 2 destination MAC, which is multicast , IPV6 do not have broadcast. Then there is a chance that you accidentally block legitimate router RA ,because there is no difference between rouge router and legitimate router that generate RA.
With private VLAN , you just add rouge router on isolated port , legitimate router with promiscuous port , everything will automatically work
upvoted 1 times

bayolo10 1 year, 8 months ago
Answer should A,https://www.geeksforgeeks.org/vlan-acl-vacl/
upvoted 2 times

pompedom 1 year, 6 months ago
It's A because PVlan limits the ability for isolated ports to communicate with other isolated ports at all, not only route advertisements.
upvoted 1 times

wts 1 year, 9 months ago
Selected Answer: D
Certain switch platforms can already implement some level of rogue RA
filtering by the administrator configuring Access Control Lists
(ACLs) that block RA ICMP messages that might be inbound on "user"
ports.

https://datatracker.ietf.org/doc/html/rfc6104#section-3.3
upvoted 1 times

steiger 2 years ago
The answer should be D
upvoted 1 times

Question #142

Topic 1

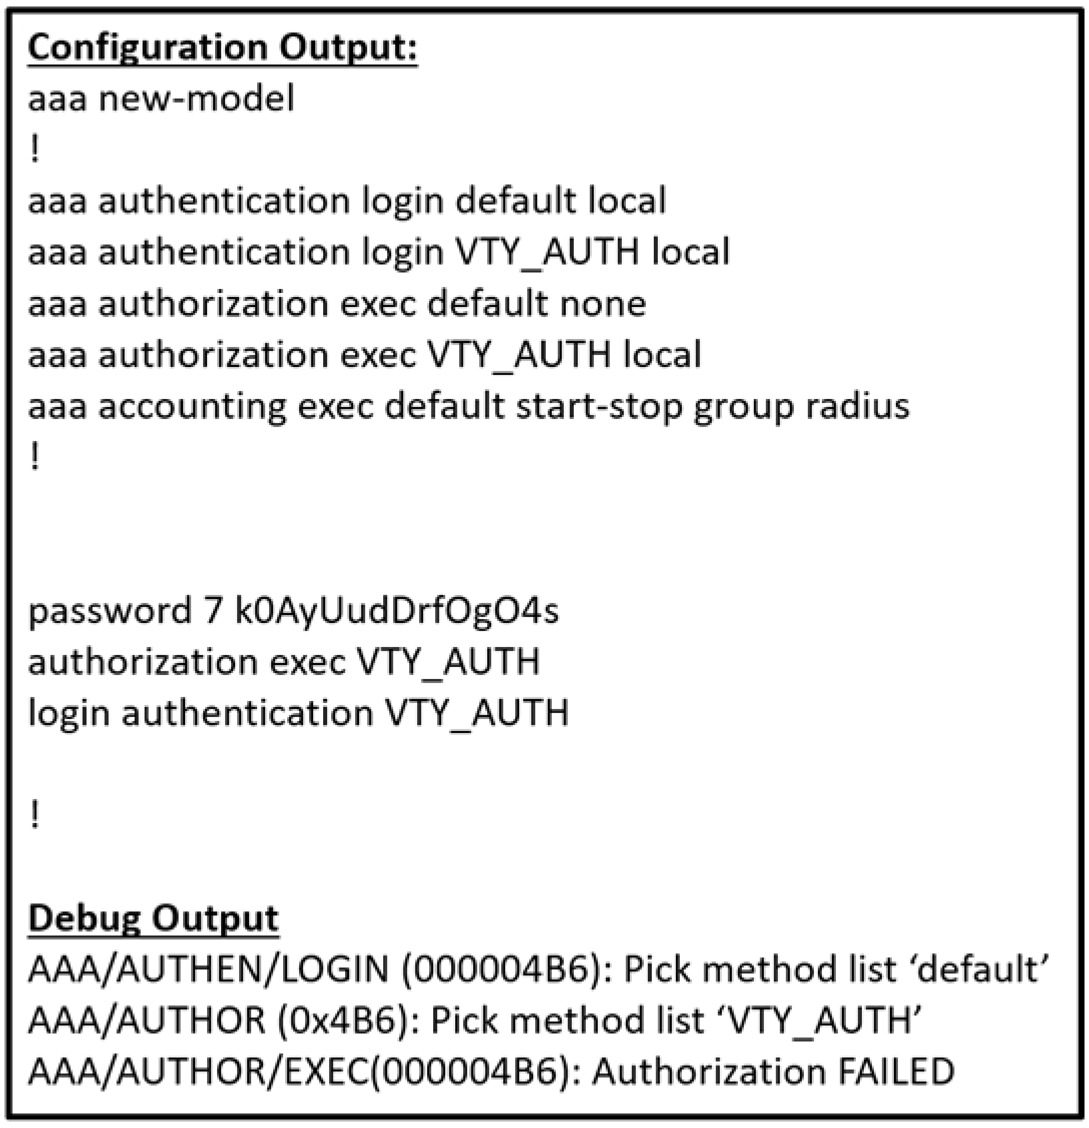

Configuration Output:

```
aaa new-model
!
aaa authentication login default local
aaa authentication login VTY_AUTH local
aaa authorization exec default none
aaa authorization exec VTY_AUTH local
aaa accounting exec default start-stop group radius
!


password 7 k0AyUudDrfOgO4s
authorization exec VTY_AUTH
login authentication VTY_AUTH

!
```

Debug Output

```
AAA/AUTHEN/LOGIN (000004B6): Pick method list 'default'
AAA/AUTHOR (0x4B6): Pick method list 'VTY_AUTH'
AAA/AUTHOR/EXEC(000004B6): Authorization FAILED
```

Refer to the exhibit.

Which action resolves the failed authentication attempt to the router?

A. Configure aaa authorization console global command

B. Configure aaa authorization console command on line vty 0 4

C. Configure aaa authorization login command on line console 0

D. Configure aaa authorization login command on line vty 0 4

Correct Answer: *A*

Reference:

https://community.cisco.com/t5/network-access-control/console-authorization-issue/td-p/2492619

Community vote distribution

A (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: A

option A:
https://community.cisco.com/t5/network-access-control/console-authorization-issue/td-p/2492619

upvoted 1 times

GreatDane 1 year, 4 months ago

Ref: Console authorization issue - Cisco Community

Post by James Horne (12-17-2015 05:37 PM)

What's missing here is the aaa authorization console command.

A. Configure aaa authorization console global command

Correct answer.

B. Configure aaa authorization console command on line vty 0 4

Wrong answer.

C. Configure aaa authorization login command on line console 0

Wrong answer.

D. Configure aaa authorization login command on line vty 0 4

Wrong answer.
upvoted 1 times

⊟ **toto2** 1 year, 9 months ago

Agree A really does nothing to fix this issue. It is a bad question with missing config information needed to actually troubleshoot this. However, the only answer that is a command that can be configured is the one shown in answer A (aaa authorization console in global config mode), so only for that reason if I would pick A. (there are "aaa authentication login" commands, but no "aaa authorization login" commands, and even the "aaa authentication login" commands are done in global config, not on the lines.) at least not on the IOS's I have seen.
upvoted 2 times

⊟ **wts** 1 year, 9 months ago

"AAA authorization is disabled on the console by default. If AAA authorization is enabled on the console, disable it by configuring the no aaa authorization console command during the AAA configuration stage. AAA should be disabled on the console for user authentication."
upvoted 1 times

⊟ **bogd** 1 year, 10 months ago

And yet if you read the full thread (https://community.cisco.com/t5/network-access-control/console-authorization-issue/td-p/2492619), the solution was NOT A...

A did nothing to fix the issue, in the end the whole AAA config on the system had to be reconfigured
upvoted 1 times

⊟ **HungarianDish** 7 months, 3 weeks ago

Yeah, still all other options are completely wrong. A) at least makes sense.
upvoted 1 times

⊟ **myrmike** 2 years ago

Debug says auth pick method was list default which implies that the user is connected to the console port. Of the answers listed only A would resolve the issue
upvoted 4 times

Question #143 *Topic 1*

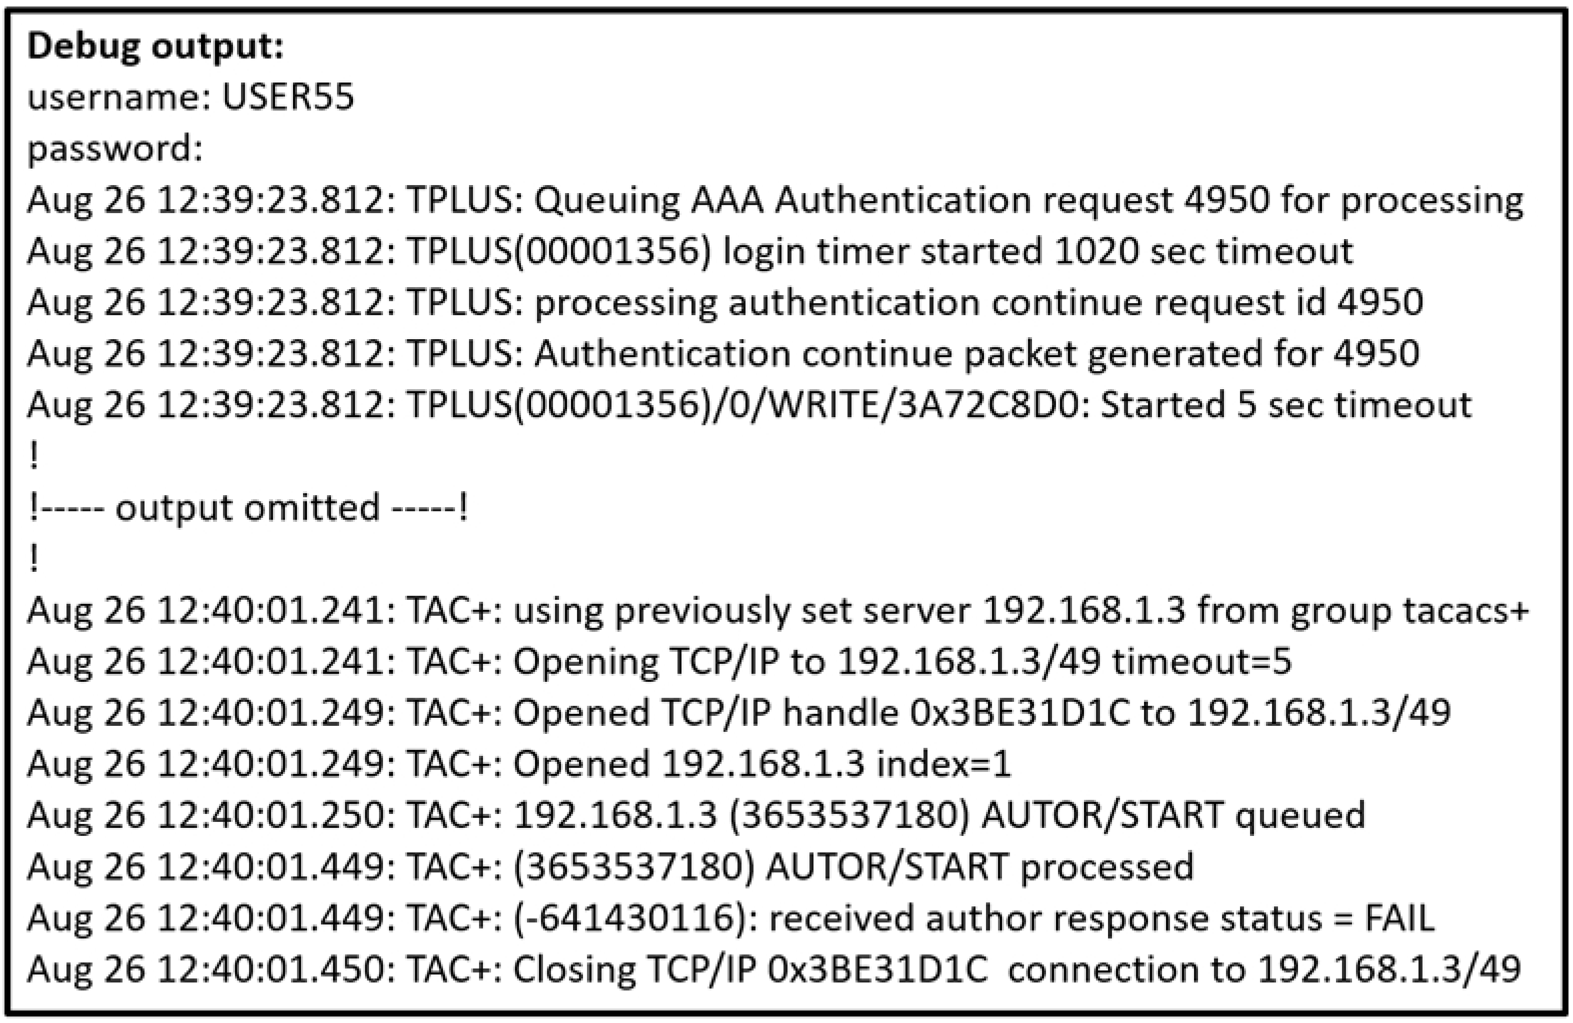

```
Debug output:
username: USER55
password:
Aug 26 12:39:23.812: TPLUS: Queuing AAA Authentication request 4950 for processing
Aug 26 12:39:23.812: TPLUS(00001356) login timer started 1020 sec timeout
Aug 26 12:39:23.812: TPLUS: processing authentication continue request id 4950
Aug 26 12:39:23.812: TPLUS: Authentication continue packet generated for 4950
Aug 26 12:39:23.812: TPLUS(00001356)/0/WRITE/3A72C8D0: Started 5 sec timeout
!
!----- output omitted -----!
!
Aug 26 12:40:01.241: TAC+: using previously set server 192.168.1.3 from group tacacs+
Aug 26 12:40:01.241: TAC+: Opening TCP/IP to 192.168.1.3/49 timeout=5
Aug 26 12:40:01.249: TAC+: Opened TCP/IP handle 0x3BE31D1C to 192.168.1.3/49
Aug 26 12:40:01.249: TAC+: Opened 192.168.1.3 index=1
Aug 26 12:40:01.250: TAC+: 192.168.1.3 (3653537180) AUTOR/START queued
Aug 26 12:40:01.449: TAC+: (3653537180) AUTOR/START processed
Aug 26 12:40:01.449: TAC+: (-641430116): received author response status = FAIL
Aug 26 12:40:01.450: TAC+: Closing TCP/IP 0x3BE31D1C  connection to 192.168.1.3/49
```

Refer to the exhibit. A network administrator logs into the router using TACACS+ username and password credentials, but the administrator cannot run any privileged commands.
Which action resolves the issue?

A. Configure the username from a local database

B. Configure TACACS+ synchronization with the Active Directory admin group

C. Configure an authorized IP address for this user to access this router

D. Configure full access for the username from TACACS+ server

Correct Answer: *D*

GreatDane 1 year, 4 months ago

Ref: TACACS+ Configuration Guide, Cisco IOS Release 15S

"CHAPTER 1
Configuring TACACS
...
How to Configure TACACS
...
Specifying TACACS Authorization

AAA authorization enables you to set parameters that restrict a user's access to the network. Authorization via TACACS+ may be applied to commands, network connections, and EXEC sessions.
..."

A. Configure the username from a local database

Wrong answer.

B. Configure TACACS+ synchronization with the Active Directory admin group

Wrong answer.

C. Configure an authorized IP address for this user to access this router

Wrong answer.

D. Configure full access for the username from TACACS+ server

Correct answer.
upvoted 1 times

Question #144

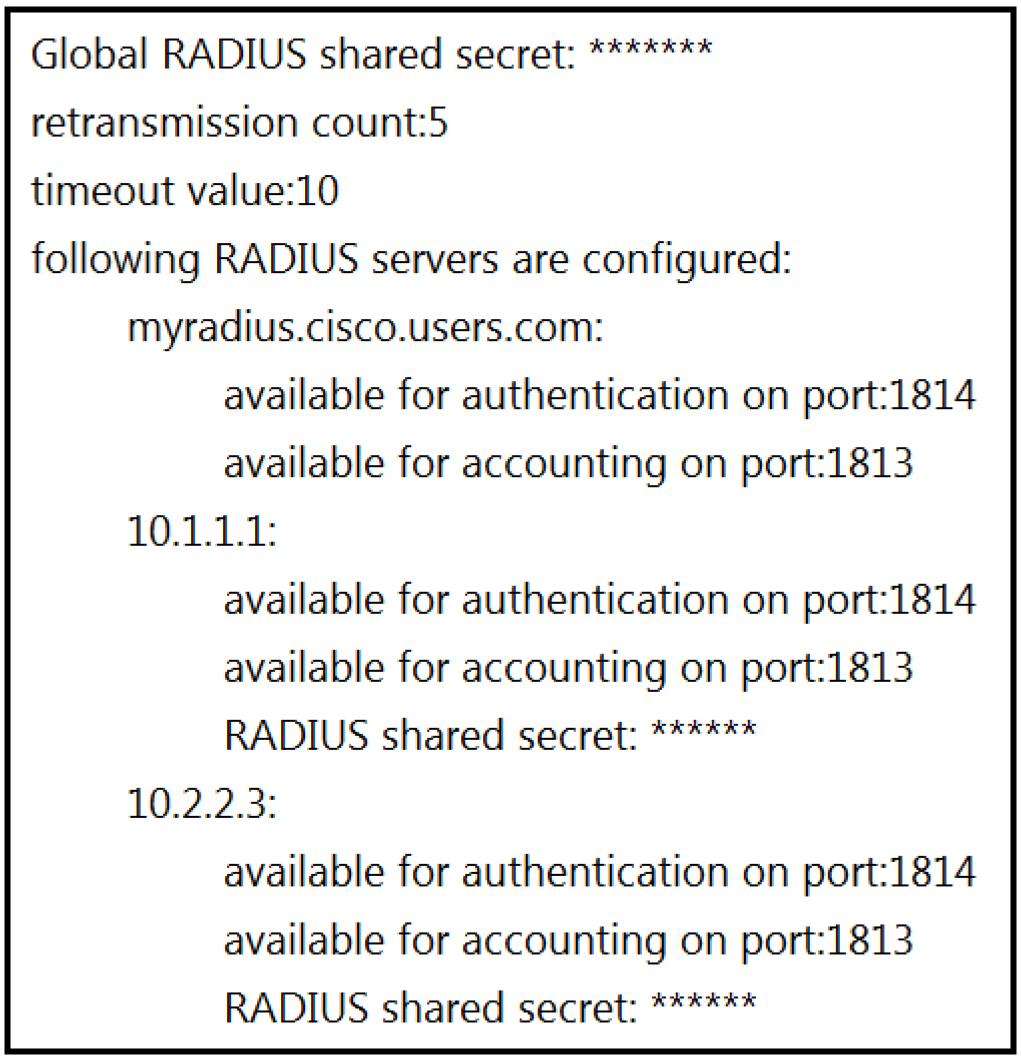

```
Global RADIUS shared secret: *******
retransmission count:5
timeout value:10
following RADIUS servers are configured:
      myradius.cisco.users.com:
            available for authentication on port:1814
            available for accounting on port:1813
      10.1.1.1:
            available for authentication on port:1814
            available for accounting on port:1813
            RADIUS shared secret: ******
      10.2.2.3:
            available for authentication on port:1814
            available for accounting on port:1813
            RADIUS shared secret: ******
```

Refer to the exhibit. AAA server 10.1.1.1 is configured with the default authentication and accounting settings, but the switch cannot communicate with the server.

Which action resolves this issue?

A. Correct the timeout value.

B. Match the authentication port.

C. Correct the shared secret.

D. Match the accounting port.

Correct Answer: *B*

⊟ **GreatDane** 1 year, 4 months ago

Ref: Solved: Which port numbers are used for RADIUS accounting and RADIUS authentication? - Cisco Community

Post by Peter Paluch

"Hi,

On all recent RADIUS server implementations, UDP/1812 is the authentication and authorization port, and UDP/1813 is the accounting port. ..."

A. Correct the timeout value.

Wrong answer.

B. Match the authentication port.

Correct answer.

C. Correct the shared secret.

Wrong answer.

D. Match the accounting port.

Wrong answer.

upvoted 3 times

⊟ **Hack4** 1 year, 10 months ago

The port values of 1812 for authentication and 1813 for accounting are RADIUS standard ports defined by the Internet Engineering Task Force (IETF) in RFCs 2865 and 2866. However, by default, many access servers use ports 1645 for authentication requests and 1646 for accounting requests.

upvoted 3 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

alalalal 2 years, 8 months ago

Radius authentication port is 1812. Hence authentication needs to be matched.

upvoted 3 times

Maurel 2 years, 8 months ago

Should be C .

upvoted 1 times

> **Dave22** 2 years, 7 months ago
>
> No its B as "default" authentication RADIUS port is 1812
>
> upvoted 4 times

Question #145 *Topic 1*

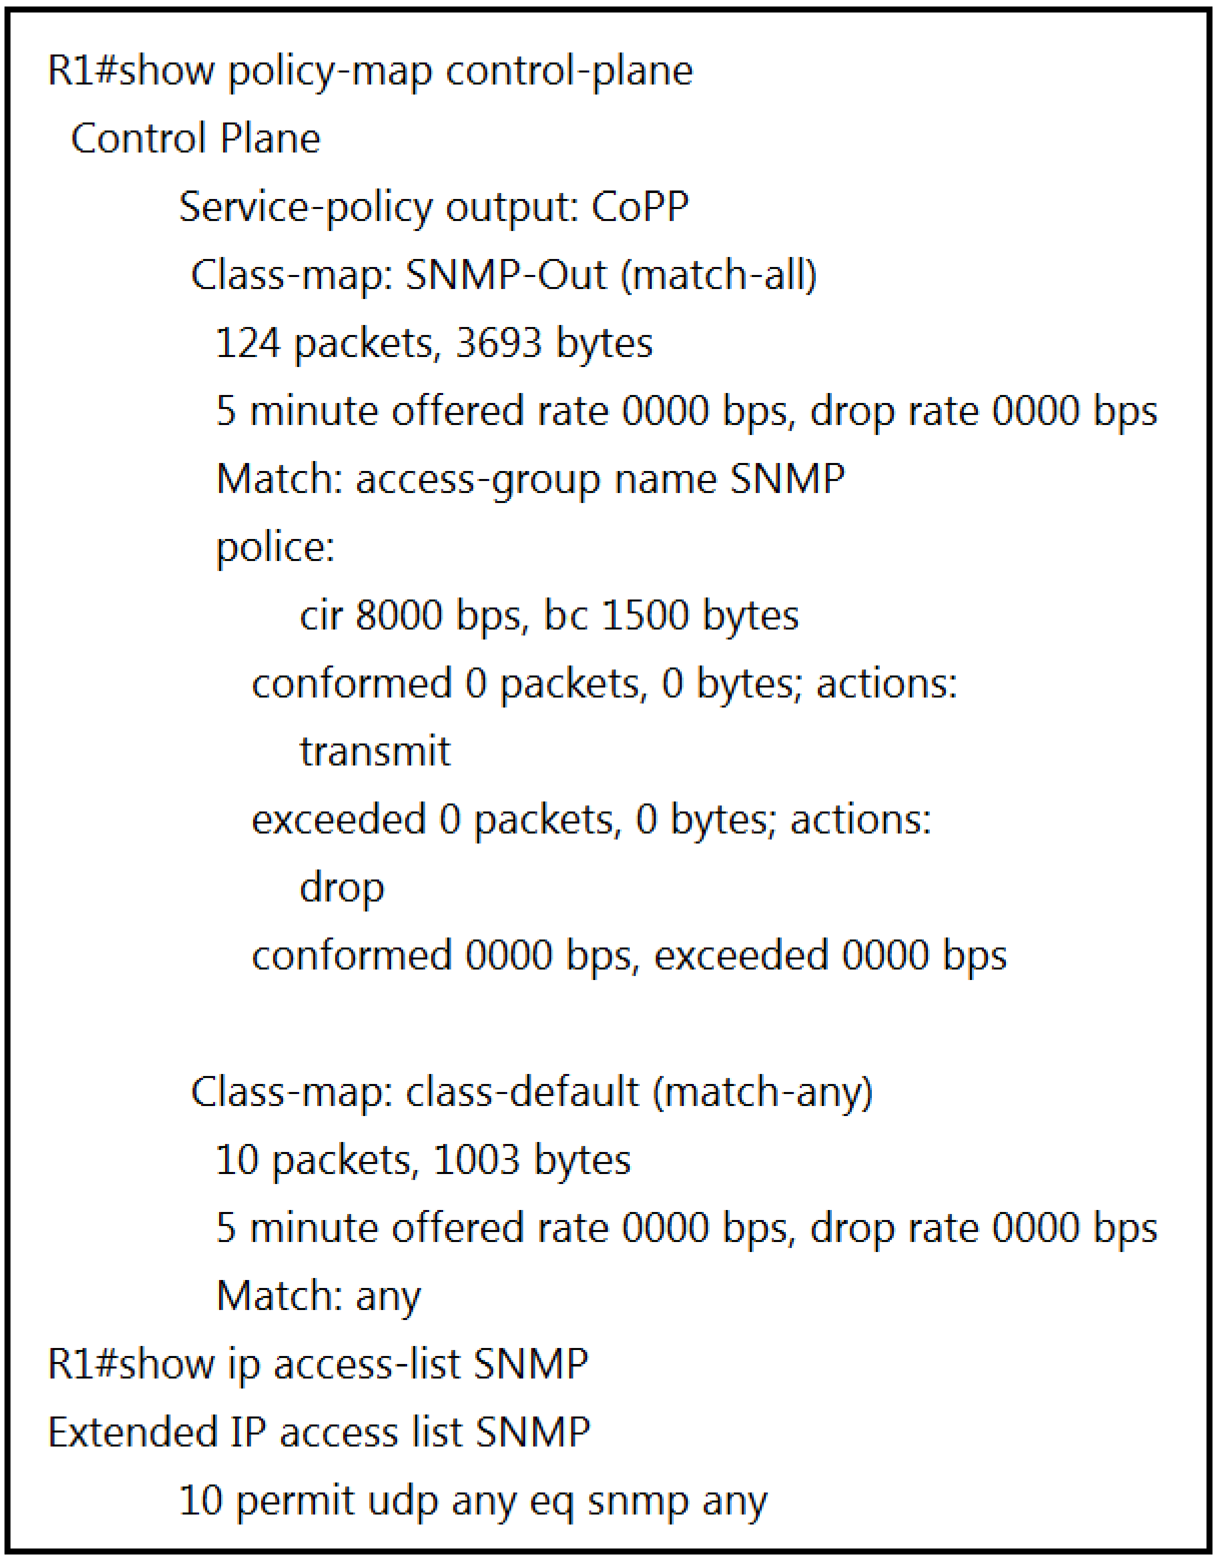

```
R1#show policy-map control-plane
  Control Plane
        Service-policy output: CoPP
         Class-map: SNMP-Out (match-all)
           124 packets, 3693 bytes
           5 minute offered rate 0000 bps, drop rate 0000 bps
           Match: access-group name SNMP
           police:
                cir 8000 bps, bc 1500 bytes
             conformed 0 packets, 0 bytes; actions:
                transmit
             exceeded 0 packets, 0 bytes; actions:
                drop
             conformed 0000 bps, exceeded 0000 bps

         Class-map: class-default (match-any)
           10 packets, 1003 bytes
           5 minute offered rate 0000 bps, drop rate 0000 bps
           Match: any
R1#show ip access-list SNMP
Extended IP access list SNMP
        10 permit udp any eq snmp any
```

Refer to the exhibit. R1 is being monitored using SNMP and monitoring devices are getting only partial information.
What action should be taken to resolve this issue?

A. Modify the CoPP policy to increase the configured exceeded limit for SNMP.

B. Modify the access list to include snmptrap.

C. Modify the CoPP policy to increase the configured CIR limit for SNMP.

D. Modify the access list to add a second line to allow udp any any eq snmp.

Correct Answer: *B*

Community vote distribution

B (88%) 13%

⊟ **Pb1805** Highly Voted 2 years, 7 months ago

The answer doesnt seem to be correct. D seems right.

Anyone?

upvoted 14 times

⊟ **Networkingguy** 1 year, 10 months ago

I think i upvoted you too soon, B seems like the better answer, tcp/ipv4 connectivity is already there. Just need to add in 162 I believe.

upvoted 3 times

⊟ **Pietjeplukgeluk** 3 weeks, 4 days ago

CoPP is applied inbound to protect your CPU from using to many cycles to process certain inbound management packets. The applied ACL on "10 permit udp any eq snmp any" is WRONG as it implies source port 161 to reach the actual router. This seems odd because the DESINATION port is actually 161 here and that one is listening on this actual router. To make the ACL actually match on inbound traffic hitting the SNMP server on this router, port 161 should be allowed as destination port as otherwise the management station cannot reach this router. Again, outbound traps should not be relavent for CoPP, if the traps overheat your CPU, it does not make a difference if they are blocked or not, the damage (high cpu) is already done. Summarazing here: the answer is D for sure as we need to allow inbound SNMP with having a destination port matching 161 == permit udep any any eq snmp (so the SNMP runs on the router, actually listening on that port) The management station is just a client in the dialog and generates a random source port.

upvoted 1 times

ytsionis Highly Voted 2 years, 1 month ago

B is the correct

snmptrap uses port 161
snmp uses port 162

ip access-list extended ABC-ACL
permit udp X.X.0.0 0.0.255.255 eq snmp host SERVER_IP !!source port is 161
permit udp X.X.0.0 0.0.255.255 host SERVER_IP eq snmptrap !!dest port is 162

https://community.cisco.com/t5/routing/acl-to-allow-snmp-traffic/td-p/1577251

upvoted 9 times

conft Most Recent 4 months, 1 week ago

Selected Answer: B

B is the correct

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

acces-lsi permit : snmp and snmptraps (agent client). The option corret is B

upvoted 1 times

adudeguy 5 months, 4 weeks ago

D
There are no matches for the traffic, so has to be related to ACL. This leaves us with B or D. The questions indicates they're getting some info and it looks like responses to SNMP requests are allowed through ACL/COPP Policy. Seems like this would just leave SNMP Traps that aren't getting out then.

upvoted 1 times

Huntkey 1 year, 2 months ago

Selected Answer: C

My apologizes... After reading the question more carefully, I would go with C. The ACL is correct. The PM is applied for outbound. So the ACL would match the response traffic from this router to the SNMP server. The class-default already matches everything so even though it is an SNMP trap, it would fall in that category and will pass. Increasing the exceed limit doesn't help because its action is to drop anyway.

upvoted 1 times

Huntkey 1 year, 2 months ago

1. Control-plane policing is only for the input direction. The question uses an "out" in the name to confuse us. The correct ACL to match SNMP poll would be in D.
SNMP trap is the output direction and it is from the router to the monitoring server so it is not affected by the control-plane policing
I would go with D

upvoted 4 times

GreatDane 1 year, 4 months ago

Device monitoring means collecting and analyzing the SNMP trap messages that devices send to the logging server. But ACL SNMP permits only SNMP traffic. This must be modified.

A. Modify the CoPP policy to increase the configured exceeded limit for SNMP.

Wrong answer.

B. Modify the access list to include snmptrap.

Correct answer.

C. Modify the CoPP policy to increase the configured CIR limit for SNMP.

Wrong answer.

D. Modify the access list to add a second line to allow udp any any eq snmp.

Wrong answer.

upvoted 2 times

Luvshah 2 months, 1 week ago

Hi, Can I have your email ID as I wanted to ask you something? Thanks.

upvoted 1 times

marcohichan 1 year, 7 months ago

B is correct. As the drop rate configured snmp is 0. Means that missing SNMP trap.

upvoted 2 times

⊟ **diogodds** 1 year, 9 months ago

In my opinion, C is the correct one, note that if SNMP traps are not included in the SNMP ACL, the CoPP class-map SNMP-Out will be skipped for that traffic, but the "class-default" will match it and will forward the traffic without policying it.

So the only viable answer is C.

upvoted 2 times

⊟ **wts** 1 year, 9 months ago

Selected Answer: B

Zeros on the counter. It seems there is no need to do something with the traffic limit.

An unspecified destination address is basically the same as "any".

Only part of the information comes to the server. Perhaps the snmp traps will complement it.

upvoted 2 times

⊟ **Hack4** 1 year, 10 months ago

"10 permit udp eq snmp any " means that : Send out only snmp informaton provide from me to any destination(mainly the NMS_SERVER). If sth like TCP event occurs in the device(SNMP_Agent as an example) is not gonna be sent to the NMS; This one is going to see only everything about UDP from the Agent . In this case to get all information provide by the Agent (R1) we need to configure snmp_trap on it....

upvoted 1 times

⊟ **Hack4** 1 year, 10 months ago

The given answer is correct. B is the right answer

upvoted 1 times

⊟ **Jenia1** 1 year, 10 months ago

My opinion is C. Modify the CoPP policy to increase the configured CIR limit for SNMP.
If you don't include the record to ACL the traffic will not be policed. so there is no reason to include Traps to the access list, and only SNMP ACL has action drop

upvoted 3 times

⊟ **JingleJangus** 1 year, 10 months ago

Selected Answer: B

id say B. Just checked the IOS and came back with this:

```
R5(config-ext-nacl)#permit udp host 2.2.2.2 eq ?
snmp Simple Network Management Protocol (161)
snmptrap SNMP Traps (162)
```

It really appears to be B, because the scenario isnt referring to intermittent access, or access to the NMS being interrupted. Its just half the picture isnt available.

upvoted 3 times

⊟ **Dirkd0344** 2 years ago

If you look closely up top the Service-policy is configured for output from the control plane virtual interface. Therefore, you would have to add an access list entry for snmptrap because the device is sending traps to the NMS.

upvoted 2 times

⊟ **Surfside92** 2 years, 1 month ago

For me D is correct answer. The access list below is wrong :
10 permit udp any eq snmp any
The allows all udp traffic from any source but it has to be from source port 161 - to any destination. Source ports vary randomly accross multiple source devices - so this would not work as an acl.

upvoted 2 times

> ⊟ **Alnet** 2 years ago
>
> The current ACL is correct. Question says this device (R1) is being monitored by another device. That means R1 is the one listening on 161. And so return packets would be sourced from R1:161. Since the ACL is applied outbound from COPP, then the source packet would be coming from source port 161. If it were applied in the inbound direction then it would be different answer.
> I think you need to add traps if only partial information is found. Traps are generally the other half of the monitoring equation in real world.
>
> upvoted 2 times
>
> ⊟ **JingleJangus** 1 year, 7 months ago
>
> https://blog.domotz.com/own-the-networks/snmp-port-number/
>
> On this page you'll note the following, from the diagram:
>
> FROM server TO agent using destination port 161 - Requests
> FROM agent TO server using destination port 161 - Responses (solicited)
> FROM agent TO server using destination port 162 - Traps (unsolicited)
>
> With SNMP, [generally] the Ephemeral Port does not apply.
>
> upvoted 1 times

⊟ **JingleJangus** 1 year, 7 months ago

Disregard my last sentence.

upvoted 1 times

Question #146 *Topic 1*

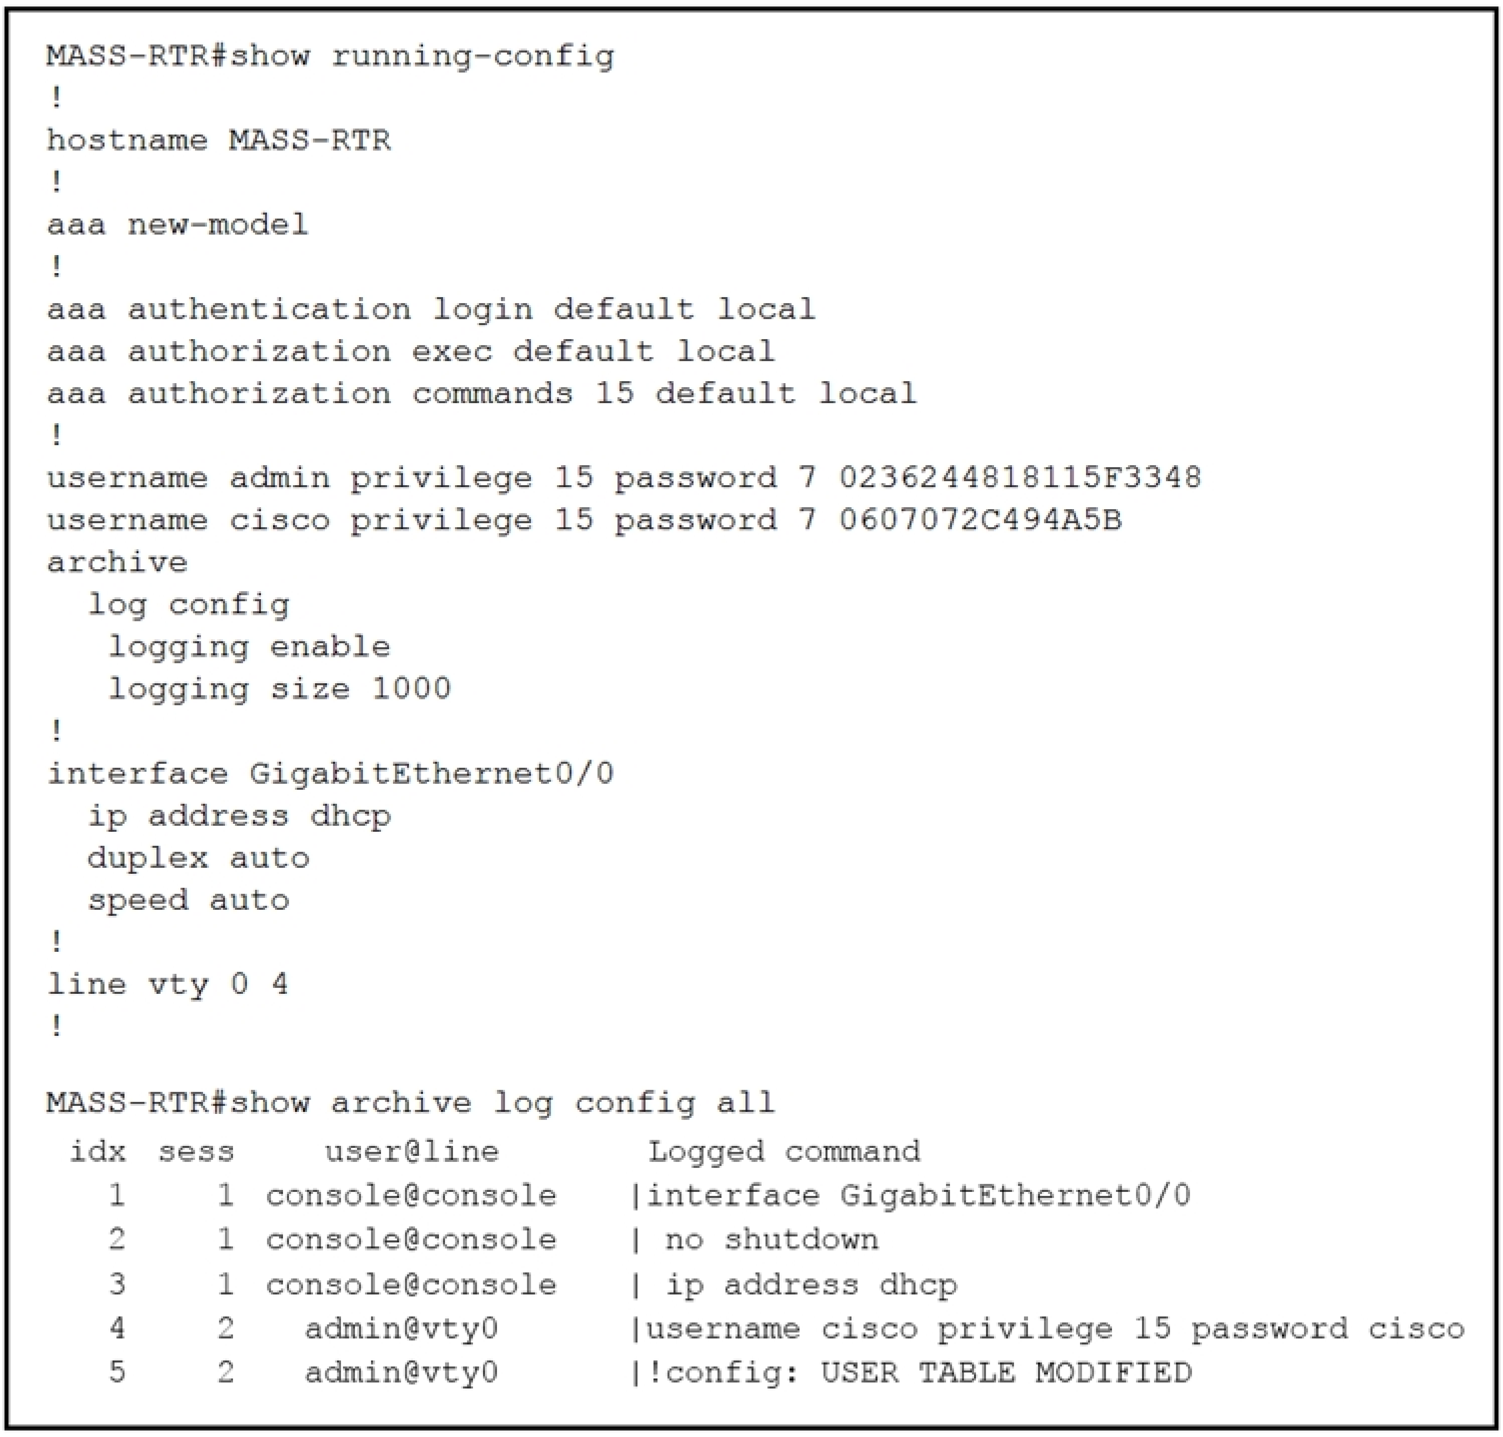

```
MASS-RTR#show running-config
!
hostname MASS-RTR
!
aaa new-model
!
aaa authentication login default local
aaa authorization exec default local
aaa authorization commands 15 default local
!
username admin privilege 15 password 7 0236244818115F3348
username cisco privilege 15 password 7 0607072C494A5B
archive
  log config
   logging enable
   logging size 1000
!
interface GigabitEthernet0/0
  ip address dhcp
  duplex auto
  speed auto
!
line vty 0 4
!

MASS-RTR#show archive log config all
 idx  sess     user@line        Logged command
   1     1  console@console    |interface GigabitEthernet0/0
   2     1  console@console    | no shutdown
   3     1  console@console    | ip address dhcp
   4     2    admin@vty0       |username cisco privilege 15 password cisco
   5     2    admin@vty0       |!config: USER TABLE MODIFIED
```

Refer to the exhibit. A client is concerned that passwords are visible when running this show archive log config all.

Which router configuration is needed to resolve this issue?

A. MASS-RTR(config)#aaa authentication arap

B. MASS-RTR(config-archive-log-cfg)#password encryption aes

C. MASS-RTR(config)#service password-encryption

D. MASS-RTR(config-archive-log-cfg)#hidekeys

Correct Answer: *D*

Community vote distribution

D (100%)

Mjestic Highly Voted 2 years, 3 months ago

Read the statement carefully. We are not talking about the "show run" (where passwords are not in plain-text) but about the "show archive log config all" (where passwords are visible).
Answer is D.

upvoted 8 times

inteldarvid Most Recent 5 months, 2 weeks ago

Selected Answer: D

D is correct:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/config-mgmt/configuration/15-sy/config-mgmt-15-sy-book/cm-config-logger.html

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

option D correct

upvoted 1 times

GreatDane 1 year, 4 months ago

Ref: Solved: Archive Command Question - Cisco Community

Post by Latchum Naidu

“Hi Pat,

Router(config-archive-log-config)# hidekeys (hides passwords from being shown / logged)
...”

A. MASS-RTR(config)#aaa authentication arap

Wrong answer.

B. MASS-RTR(config-archive-log-cfg)#password encryption aes

Wrong answer.

C. MASS-RTR(config)#service password-encryption

Wrong answer.

D. MASS-RTR(config-archive-log-cfg)#hidekeys

Correct answer.

upvoted 1 times

toni2 1 year, 10 months ago

Correct Answer D
Last but not least, it might be a good idea not to store any passwords in the configuration change logs. You can use the following command to disable this:
Router(config-archive-log-cfg)#hidekeys

https://networklessons.com/cisco/ccie-routing-switching/configuration-change-notification-logging

upvoted 1 times

leecharxos 1 year, 11 months ago

Totally agree : (Optional) Suppresses the display of password information in configuration log files.Enabling the "hidekeys command" increases security by preventing password information from being displayed in configuration log files.....https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/config-mgmt/configuration/15-sy/config-mgmt-15-sy-book/cm-config-logger.pdf

upvoted 1 times

irukmana97 2 years, 2 months ago

Tested on the lab, the given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Dejjie 10 months ago

All answers are always right to you.

upvoted 1 times

Hodepine77 2 years, 5 months ago

Tested this on some live equipment, it's the hidekeys command.

upvoted 1 times

puggy88 2 years, 6 months ago

i think its C

upvoted 1 times

Question #147 Topic 1

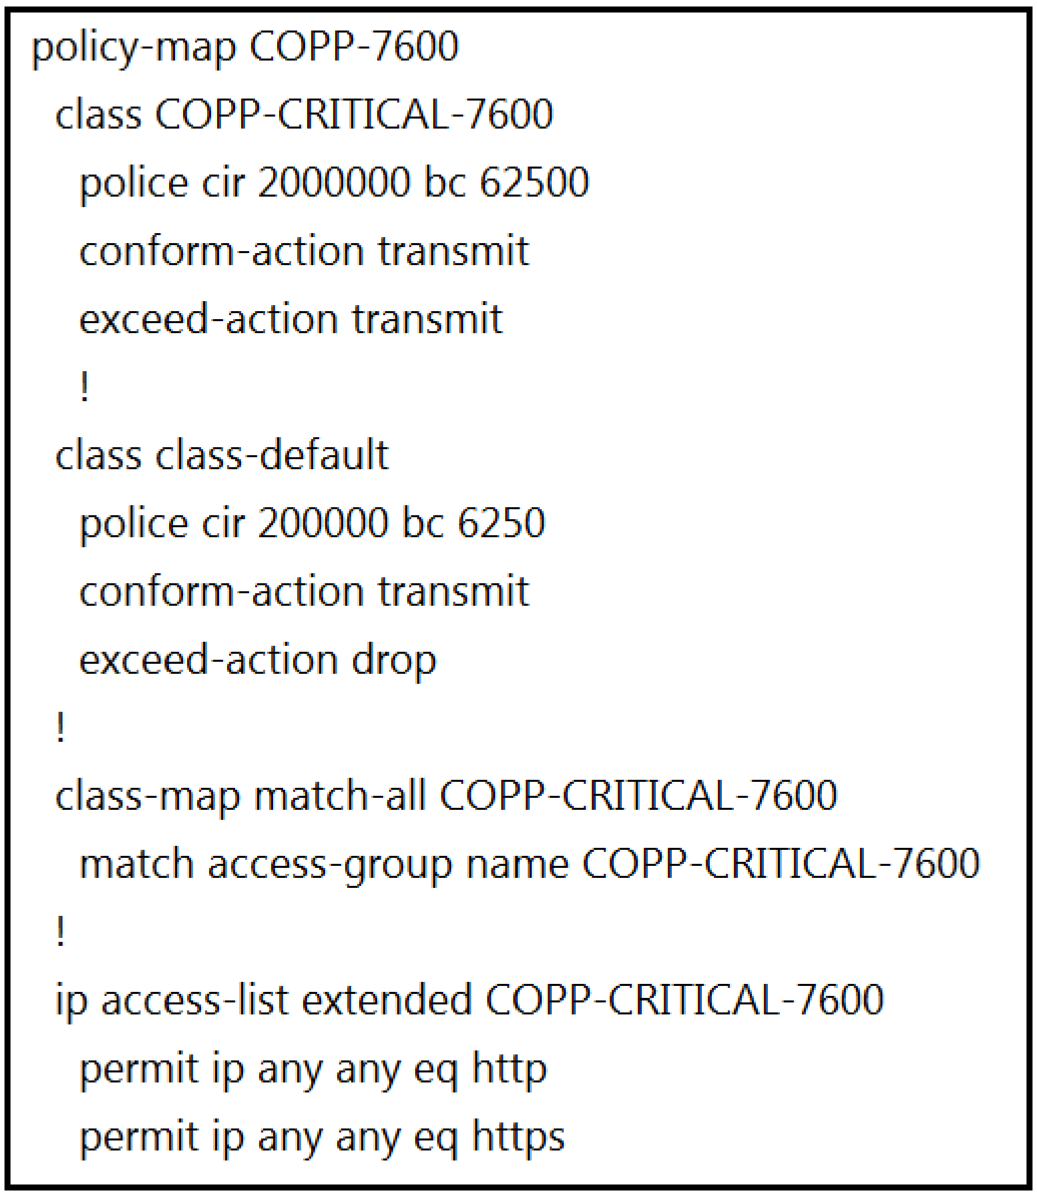

```
policy-map COPP-7600
  class COPP-CRITICAL-7600
    police cir 2000000 bc 62500
    conform-action transmit
    exceed-action transmit
    !
  class class-default
    police cir 200000 bc 6250
    conform-action transmit
    exceed-action drop
  !
  class-map match-all COPP-CRITICAL-7600
    match access-group name COPP-CRITICAL-7600
  !
  ip access-list extended COPP-CRITICAL-7600
    permit ip any any eq http
    permit ip any any eq https
```

Refer to the exhibit. BGP is flapping after the CoPP policy is applied.
What are the two solutions to fix the issue? (Choose two.)

A. Configure a higher value for CIR under the Class COPP-CRITICAL-7600.

B. Configure a higher value for CIR under the default class to allow more packets during peak traffic.

C. Configure BGP in the COPP-CRITICAL-7600 ACL.

D. Configure IP CEF for CoPP policy and BGP to work.

E. Configure a three-color policer instead of two-color policer under Class COPP-CRITICAL-7600.

Correct Answer: *BC*

Community vote distribution

BC (86%) 14%

⊟ **inteldarvid** 5 months, 2 weeks ago

Selected Answer: BC

B and C:
Explanation/Reference:
Explanation:
The policy-map COPP-7600 only rate-limit HTTP & HTTPS traffic (based on the ACLconditions) so any BGP packets will be processed in the class "class-default", which dropsexceeded BGP packets. Therefore we have two ways to solvethis problem:
+ Add BGP to the ACL with the statement "permit tcp any any eq bgp"
+ Configure higher value for CIR in default class as 2Mbps is too low for web traffic (http & https)

upvoted 1 times

⊟ **Huntkey** 1 year, 2 months ago

Selected Answer: BC

C takes care when the BGP session is initiated from the peer router
B takes care when the BGP session is initialized from the local router. In this case, the traffic coming in would have destination port of a random number. It would match the default class.

upvoted 3 times

⊟ **Remsync** 1 year, 2 months ago

Selected Answer: BC

B and C are correct.
upvoted 2 times

pompedom 1 year, 6 months ago

Selected Answer: AC

You have to increase cir of copp critical not the default one. remember bgp is part of COPP-CRITICAL-7600 now
upvoted 1 times

sajjad_gayyem 6 months, 1 week ago

But the exceed action is transmit in A.
upvoted 1 times

JingleJangus 1 year, 6 months ago

Not Necessary.

COPP-CRITICAL-7600 is configured as;
confirm - transmit
exceed - transmit
Meaning traffic is never dropped, regardless of how high, or low, the CIR is configured as.

Question is asking for 2 different solutions, NOT 2 elements of the same solution.

If the engineer does not want to add BGP to COPP-CRITICAL-7600, another solution is to increase the CIR of class-default, so as to reduce the chances that traffic is dropped, including BGP.
upvoted 10 times

wts 1 year, 9 months ago

Why change the default settings if bgp falls into COPP-CRITICAL-7600?
For bgp and http(s) you need to make different policies. But I don't see such an option.
upvoted 3 times

Hack4 1 year, 10 months ago

The given answer is correct
upvoted 2 times

error_909 2 years, 3 months ago

The given answer is correct
upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct
upvoted 2 times

Question #148 Topic 1

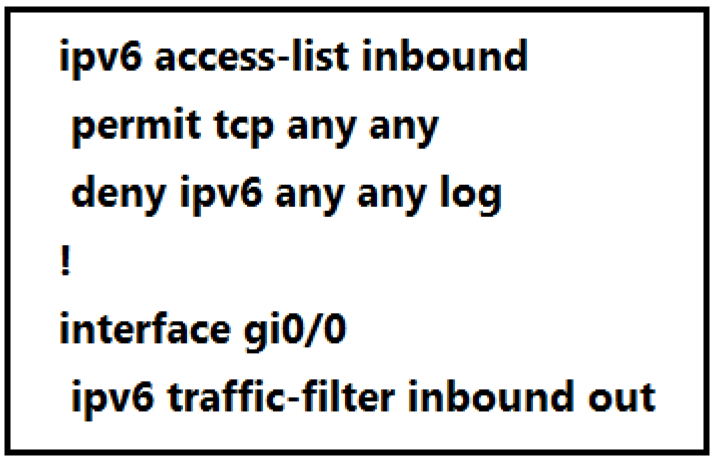

```
ipv6 access-list inbound
 permit tcp any any
 deny ipv6 any any log
!
interface gi0/0
 ipv6 traffic-filter inbound out
```

Refer to the exhibit. A network administrator configured an IPv6 access list to allow TCP return traffic only, but it is not working as expected. Which changes resolve this issue?

A.

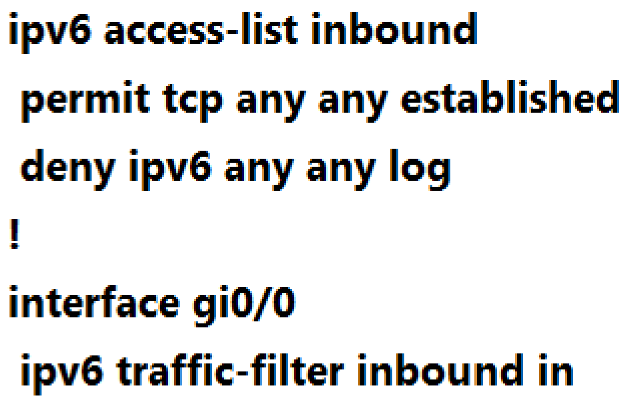

```
ipv6 access-list inbound
 permit tcp any any established
 deny ipv6 any any log
!
interface gi0/0
 ipv6 traffic-filter inbound in
```

B.

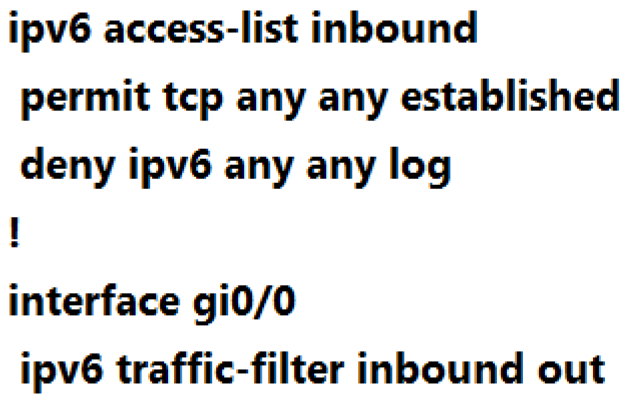

```
ipv6 access-list inbound
 permit tcp any any established
 deny ipv6 any any log
!
interface gi0/0
 ipv6 traffic-filter inbound out
```

C.

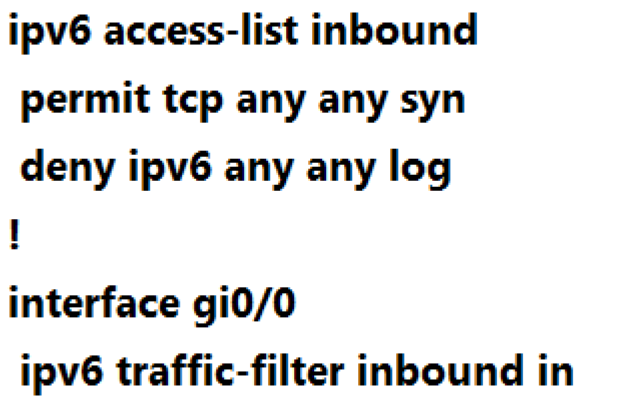

```
ipv6 access-list inbound
 permit tcp any any syn
 deny ipv6 any any log
!
interface gi0/0
 ipv6 traffic-filter inbound in
```

D.

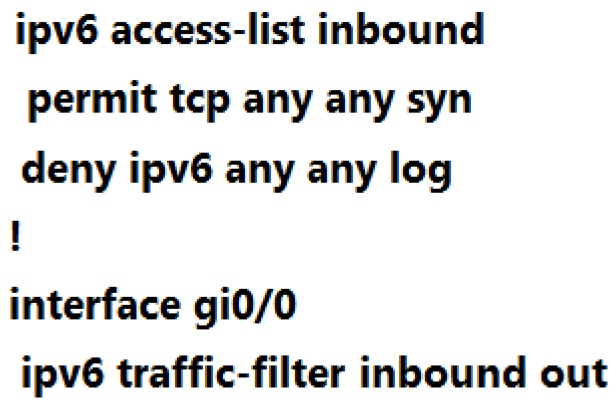

```
ipv6 access-list inbound
 permit tcp any any syn
 deny ipv6 any any log
!
interface gi0/0
 ipv6 traffic-filter inbound out
```

Correct Answer: *A*

⊟ **GreatDane** 1 year, 4 months ago

TCP hosts establish a connection-oriented session with one another using a “three-way handshake” mechanism.
As far as I know, the TCP return frame is the last frame involved in the three-way handshake (the ACK frame). Then, the session between the two hosts is established.

So:

permit tcp any any established (let the TCP return frame in, from any host)
deny ipv6 any any log (deny any other IPv6 traffic from any host)

Since the TCP return frame must be allowed IN, the ACL must be applied IN.

Answer A is correct.

upvoted 4 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 3 times

Question #149 *Topic 1*

What are two functions of IPv6 Source Guard? (Choose two.)

A. It works independent from IPv6 neighbor discovery.

B. It denies traffic from unknown sources or unallocated addresses.

C. It uses the populated binding table to allow legitimate traffic.

D. It denies traffic by inspecting neighbor discovery packets for specific patterns.

E. It blocks certain traffic by inspecting DHCP packets for specific sources.

Correct Answer: *BC*

Community vote distribution

BC (67%) A (33%)

chris110 3 months, 3 weeks ago

Selected Answer: BC

B. It denies traffic from unknown sources or unallocated addresses.
C. It uses the populated binding table to allow legitimate traffic.

upvoted 1 times

Brand 3 months, 3 weeks ago

Selected Answer: BC

First of all the question asks to choose two. Second of all, as the name indicates the Source Guard feature determines if the source of a traffic is coming from a prefix or address in the binding table. Binding table entries are populated using mechanisms like ND. So saying "It works independent from IPv6 neighbor discovery." is WRONG. So "one" of the two correct answers can not be A.

upvoted 1 times

conft 4 months, 1 week ago

Selected Answer: A

the given answer is correct.

upvoted 1 times

GreatDane 1 year, 4 months ago

Ref: IPv6 Source Guard and Prefix Guard – Cisco

"...
Information About IPv6 Source Guard and Prefix Guard

IPv6 Source Guard Overview

IPv6 source guard is an interface feature between the populated binding table and data traffic filtering. This feature enables the device to deny traffic when it is originated from an address that is not stored in the binding table.
...
IPv6 source guard can deny traffic from unknown sources or unallocated addresses, such as traffic from sources not assigned by a DHCP server.
..."

A. It works independent from IPv6 neighbor discovery.

Wrong answer.

B. It denies traffic from unknown sources or unallocated addresses.

Correct answer.

C. It uses the populated binding table to allow legitimate traffic.

Correct answer.

D. It denies traffic by inspecting neighbor discovery packets for specific patterns.

Wrong answer.

E. It blocks certain traffic by inspecting DHCP packets for specific sources.

Wrong answer.

upvoted 3 times

examShark 2 years, 4 months ago

The given answer is correct
IPv6 Source Guard blocks any data traffic from an unknown source. For example, one that is not already populated in the binding table or previously learned through Neighbor Discovery (ND) or Dynamic Host Configuration Protocol (DHCP) gleaning.

upvoted 3 times

leecharxos 1 year, 11 months ago

yeap :It filters inbound traffic on L2 switch ports that are not in the IPv6 binding table,
https://networklessons.com/cisco/ccie-routing-switching-written/ipv6-source-guard

upvoted 2 times

Question #150 Topic 1

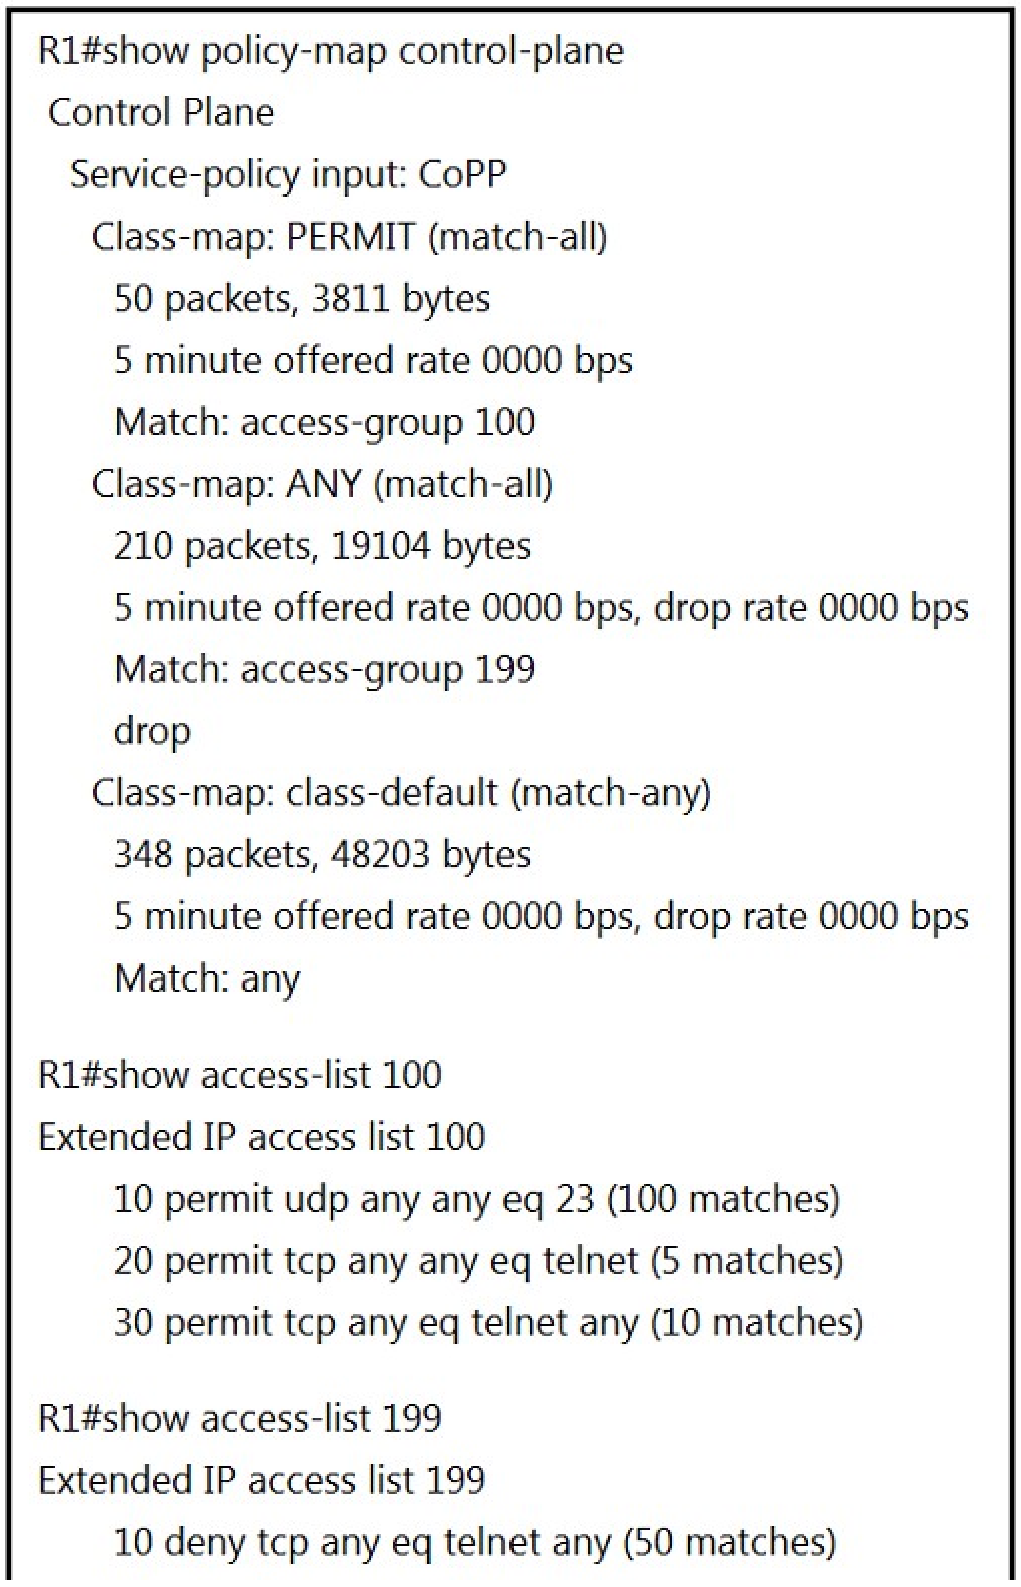

```
R1#show policy-map control-plane
 Control Plane
   Service-policy input: CoPP
     Class-map: PERMIT (match-all)
       50 packets, 3811 bytes
       5 minute offered rate 0000 bps
       Match: access-group 100
     Class-map: ANY (match-all)
       210 packets, 19104 bytes
       5 minute offered rate 0000 bps, drop rate 0000 bps
       Match: access-group 199
       drop
     Class-map: class-default (match-any)
       348 packets, 48203 bytes
       5 minute offered rate 0000 bps, drop rate 0000 bps
       Match: any

R1#show access-list 100
Extended IP access list 100
       10 permit udp any any eq 23 (100 matches)
       20 permit tcp any any eq telnet (5 matches)
       30 permit tcp any eq telnet any (10 matches)

R1#show access-list 199
Extended IP access list 199
       10 deny tcp any eq telnet any (50 matches)
```

Refer to the exhibit. Which two actions restrict access to router R1 by SSH? (Choose two.)

A. Remove class-map ANY from service-policy CoPP.

B. Configure transport output ssh on line vty and remove sequence 20 from access list 100.

C. Configure transport input ssh on line vty and remove sequence 30 from access list 100.

D. Remove sequence 10 from access list 100 and add sequence 20 deny tcp any any eq telnet to access list 199.

E. Configure transport output ssh on line vty and remove sequence 10 from access list 199.

Correct Answer: *AC*

Community vote distribution

AC (81%) BC (19%)

DaanB Highly Voted 2 years, 8 months ago

B and C. A is not correct - IMO

upvoted 11 times

bjromero28 Highly Voted 2 years, 1 month ago

This image is cut off. Here's the is continuation below:

R1# show access-list 199
Extended ip access list 199
10 deny tcp any eq telnet any (50 matches)
50 permit ip any any (1 match)

R1# show running-config | section line vty
line vty 0 4

```
login
transport input telnet ssh
transport output telnet ssh
------------------------------------------------------------------------
```

In order to restrict access to ssh only, shouldn't we limit the vty lines to transport ssh only?

I believe the answer is B and C.

upvoted 9 times

spapi0390 2 years ago

I have done that on lab, with the above output the SSH is not working! So i have remove Class-map ANY- then I was able to SSH to the router. So A is 100% ok. Other best option is C, since if we replace input telnet ssh to only SSH then you do not have access through telnet on the router.

upvoted 4 times

AlexInShort12 Most Recent 3 days, 12 hours ago

Not clear question, not sure if we are suppose to
allow connection GOINGTO R1 via SSH
or
Allow R1 making SSH connection out only via SSH.

upvoted 1 times

net_eng10021 3 months ago

Awfully worded question....

upvoted 1 times

conft 4 months, 1 week ago

Selected Answer: AC

A and C is the correct.

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: AC

AC is correct

upvoted 1 times

Malasxd 7 months, 1 week ago

Selected Answer: AC

A and C are right.
A) ACL 199 match SSH traffic by sequence 50. The class-map match ACL 199 and this class is droping all traffic. if you remove the SSH traffic will match default class and will pass. If you don't permit SSH in ACL 100 it's mandatory remove this class.

B) if you configure output ssh you are allowing R1 being the connection's client and i'm not sure if it is desided by the question. but you need to configure SSH input to ssh works and there is no option to do it except option C.

C) It works with option A. Mandatory you need to input ssh in the lines vty to allow SSH and this is the unique option you can do it. We don't have the option to include SSH in ACL 100, so we need to remove the class ANY and input the SSH. Option C also removes sequence 30 in ACL 100 and this make the router unable to answer telnet connection. I would prefer to remover sequence 20, but removing sequence 30 also works.

D) Does not make sense to me.

E) does not make sense either.

upvoted 3 times

Clarent_I 5 months, 2 weeks ago

Removing Sequence 30 in AC doesn't make the router unable to answer telnet connection. It is simply disallowing the remote device to respond back to the connection initiated by R1 because the control plane has the service policy applied in inbound direction. Hence Option B is not needed to be used to stop the outbound SSH connection thou the question never asked for this.
Thou, your explanations for A and C being the right answers are correct.

upvoted 2 times

Pietjeplukgeluk 3 weeks, 3 days ago

I think this question is wrong as removing class ANY will mean you do not use CoPP at all. If the technology provides any benefits, why have questions that just allow all traffic? Anyway, i would not mind making a question like this wrong.

upvoted 1 times

ericxw 11 months, 3 weeks ago

Selected Answer: AC

transport output ssh --- this will allow only ssh to be initiated from this device - which is not required - so A & C

upvoted 1 times

NoUserName1234 1 year ago

Selected Answer: BC

Full picture seen on the following site givin picture is wrong.

https://www.actual4test.com/articles/dec-2021-pass-300-410-exam-in-first-attempt-updated300-410-actual4test-exam-question-q91-q113/

upvoted 2 times

⊟ **Huntkey** 1 year, 2 months ago

Class ANY will match pretty much everything. The only thing it doesn't match is the outbound telnet from the router to where else (because the seq 10 in ACL 199 would match the return traffic). Therefore, you must remove this class because it would deny the inbound SSH traffic
C would restrict inbound to be SSH only, despite that the "PERMIT" map would allow for inbound Telnet

upvoted 2 times

⊟ **wts** 1 year, 3 months ago

Selected Answer: AC

It seems that it is necessary to reduce the options for connecting to the router to SSH.
Block telnet, allow SSH - it's clearer.

Only the ANY captures(ACL199) SSH packets for policy(only this class-map can influence the ssh by control plane policy):
10 deny tcp any eq telnet any
50 permit ip any any <--------------------here(picture cropped)
i.e. A

By removing the ANY, we will skip the ssh packages default class. But apparently, "restrict" means that you need to disable telnet, leaving only ssh TO router.
So we need the command "transport input ssh",
i.e. C.

P.S.: disgusting question

upvoted 2 times

⊟ **TECH3K3** 1 year, 5 months ago

B and C
Some configuration output is missing, which is why some of you are choosing the wrong answers. See below for missing VTY Line config.
line vty 0 4
transport input telnet ssh
transport output telnet ssh
We only want SSH and no Telnet session.
Configuring transport input/output ssh with remove the transport input telnet off the vty line.
Also if you select B and C, you will also remove telnet from ACL 100.

upvoted 1 times

⊟ **TECH3K3** 1 year, 5 months ago

Selected Answer: BC

B and C
We only want SSH and no Telnet session.
Configuring transport input/output ssh with remove the transport input telnet off the vty line.
Also if you select B and C, you will also remove telnet from ACL 100.

upvoted 1 times

⊟ **Carl1999** 1 year, 10 months ago

Selected Answer: AC

I understood the meaning of the sentence, it means that ONLY SSH CAN CONNECT.
A and C.

upvoted 2 times

⊟ **Carl1999** 1 year, 10 months ago

I think the following is easier.
access list 100
40 permit tcp any any eq 22.

upvoted 1 times

⊟ **wts** 1 year, 11 months ago

I don't see it having anything to do with blocking access via ssh.

upvoted 3 times

⊟ **OhBee** 1 year, 11 months ago

Selected Answer: AC

A is correct. Note that once the ANY class-map is removed, SSH traffic will match the default class-map, which transmits all remaining traffic.

upvoted 4 times

⊟ **MarvinY** 1 year, 11 months ago

Selected Answer: AC

A. SSH traffic is matching line 50 of the class-map ANY and getting dropped. So class-map ANY needs to be removed to allow the SSH connection
B. I believe this is for outbound connection, not relevant to the question.
C. Correct.

upvoted 3 times

⊟ **MarvinY** 1 year, 11 months ago

A. SSH traffic is matching line 50 of the ACL 199 and getting dropped by class-map ANY. So class-map ANY needs to be removed to allow the SSH connection

upvoted 2 times

Question #151 Topic 1

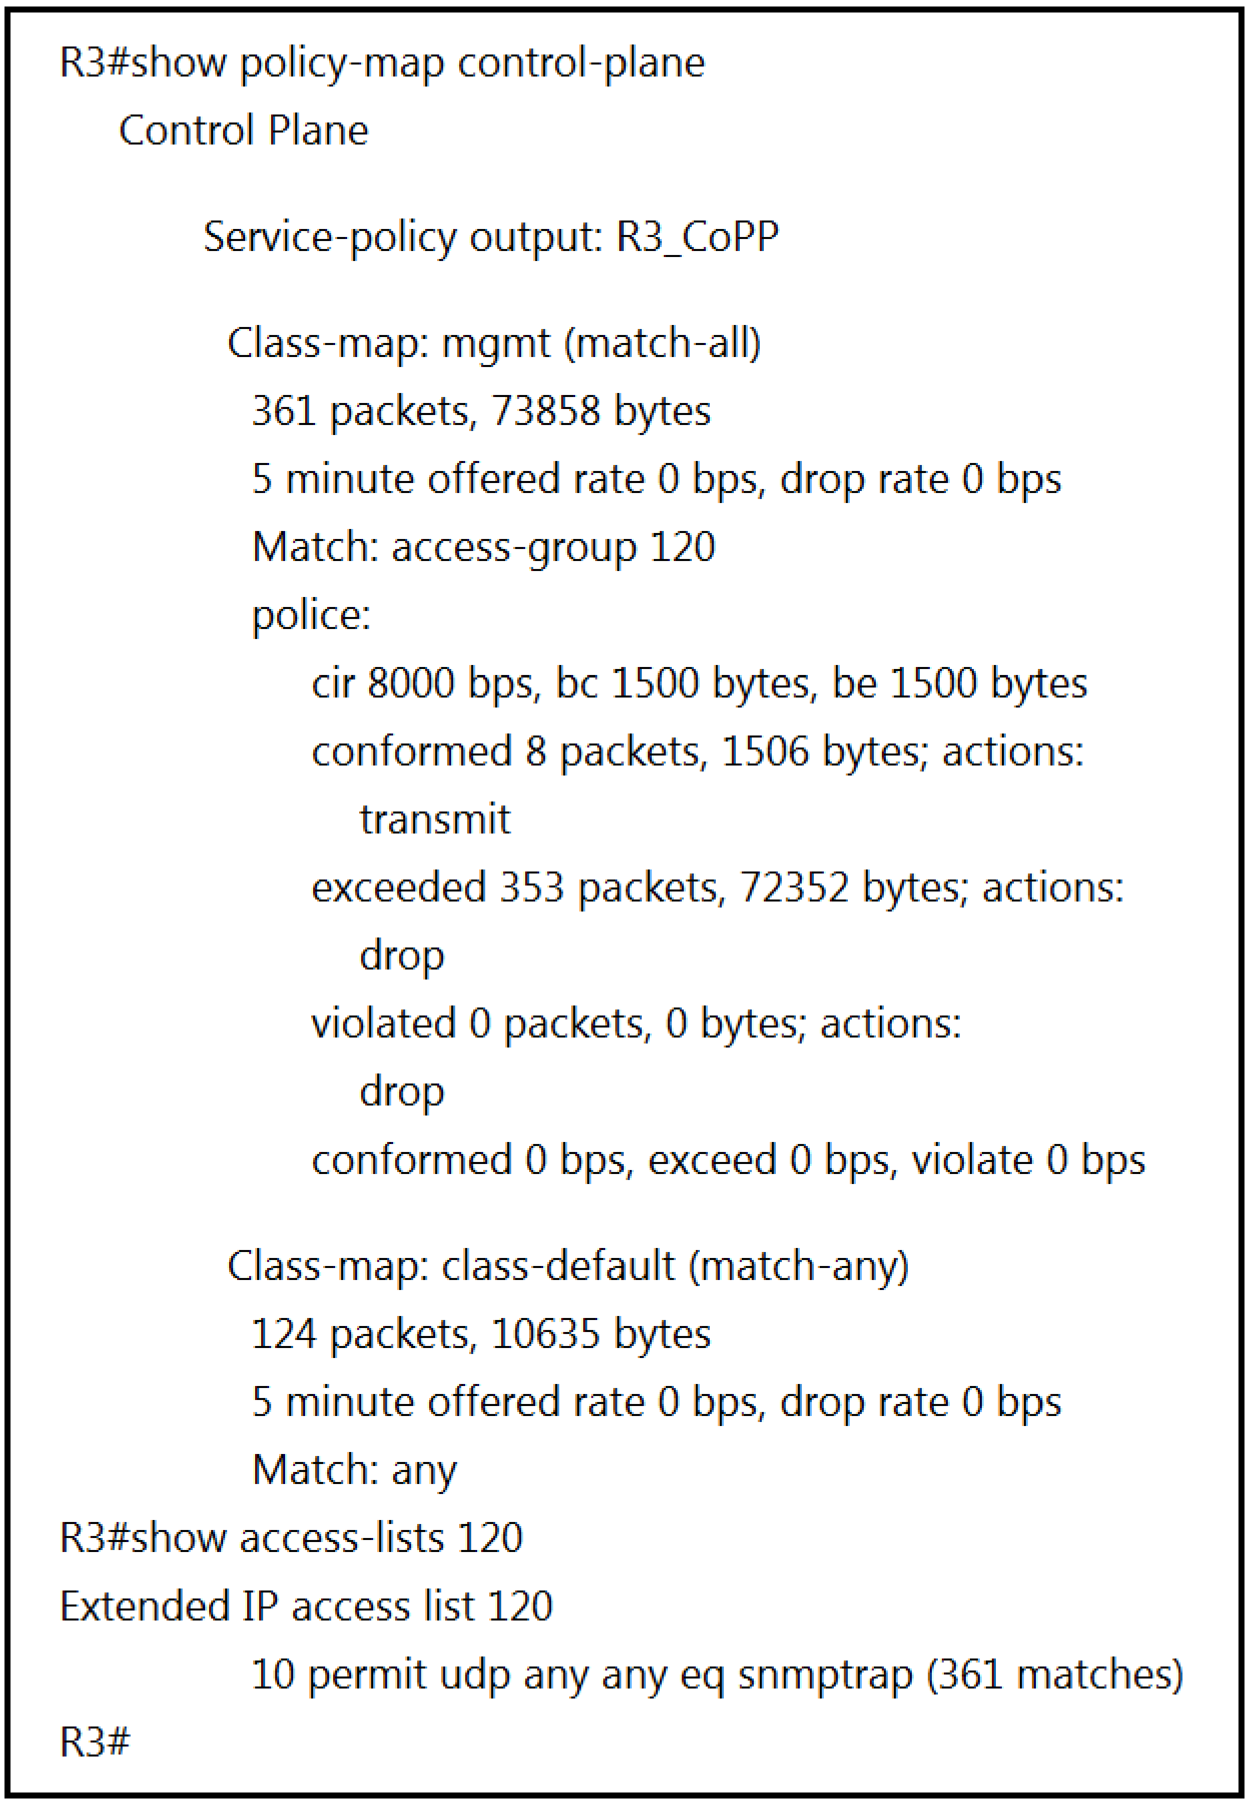

```
R3#show policy-map control-plane
   Control Plane

       Service-policy output: R3_CoPP

         Class-map: mgmt (match-all)
           361 packets, 73858 bytes
           5 minute offered rate 0 bps, drop rate 0 bps
           Match: access-group 120
           police:
               cir 8000 bps, bc 1500 bytes, be 1500 bytes
               conformed 8 packets, 1506 bytes; actions:
                  transmit
               exceeded 353 packets, 72352 bytes; actions:
                  drop
               violated 0 packets, 0 bytes; actions:
                  drop
               conformed 0 bps, exceed 0 bps, violate 0 bps

         Class-map: class-default (match-any)
           124 packets, 10635 bytes
           5 minute offered rate 0 bps, drop rate 0 bps
           Match: any
R3#show access-lists 120
Extended IP access list 120
           10 permit udp any any eq snmptrap (361 matches)
R3#
```

Refer to the exhibit. Which action resolves intermittent connectivity observed with the SNMP trap rackets?

A. Decrease the committed burst size of the mgmt class map.

B. Increase the CIR of the mgmt class map.

C. Add one new entry in the ACL 120 to permit the UDP port 161.

D. Add a new class map to match TCP traffic.

Correct Answer: *B*

Community vote distribution

B (100%)

Malasxd 7 months, 1 week ago

Selected Answer: B

B is right.
The firs line says "Service-policy OUTPUT", so the traffic is originated by the router. We also can see the class mgmt has exceeded matches and it is dropping packets.

A) does not make sense. the class already dropping packet, if you decrease the CIR it will get even worse.

B) is right. The class is dropping the packets because they exceeded the bandwidth. If you increase to the correct bandwidth it will work.

C) SNMTP traps works in port 162.

D) does not make sense
upvoted 2 times

⊟ **Huntkey** 1 year, 2 months ago

Control-plane policing matches the inbound traffic. SNMPTRAP are outbound traffic. the policy doesn't affect it at all. The match in the ACL would mean other devices sending traps to this local router or something. SNMPTRAP is connection less and stateless. There is no such thing for intermittent connectivity for SNMPTRAP traffic.

My understanding is that some TCP connection got disconnected constantly because of the class-map and the router is sending out SNMPTRAP to notify people about it. Therefore, I would go with D
upvoted 1 times

⊟ **Huntkey** 1 year, 2 months ago

Never mind... Apparently the policy-map can also be applied for the output direction... B is correct
upvoted 1 times

⊟ **GreatDane** 1 year, 4 months ago

As the exhibit shows, among all matches by ACL 120 (361 packets), 353 exceeded the CIR of class map mgmt, while only 8 packets conformed to it.
Here, the first thing to do is to give more bandwidth to class map mgmt.

A. Decrease the committed burst size of the mgmt class map.

Wrong answer.

B. Increase the CIR of the mgmt class map.

Correct answer.

C. Add one new entry in the ACL 120 to permit the UDP port 161.

Wrong answer.

D. Add a new class map to match TCP traffic.

Wrong answer.
upvoted 2 times

⊟ **error_909** 2 years, 3 months ago

The given answer is correct
upvoted 2 times

⊟ **examShark** 2 years, 4 months ago

The given answer is correct
upvoted 3 times

⊟ **Vince64** 2 years, 6 months ago

Connectivity is intermittent so C and D may not be the correct answer
upvoted 1 times

⊟ **willlee** 2 years, 7 months ago

anybody?

i think its C
upvoted 3 times

⊟ **spapi0390** 2 years ago

it could be but as far as exceeded actions its dropping then the given answer is correct
upvoted 1 times

⊟ **Alnet** 2 years ago

No. SNMP Trap is on port 162. Port 161 is just for SNMP requests (get, inform...), but not SNMP Traps.
upvoted 2 times

Question #152 Topic 1

DRAG DROP -

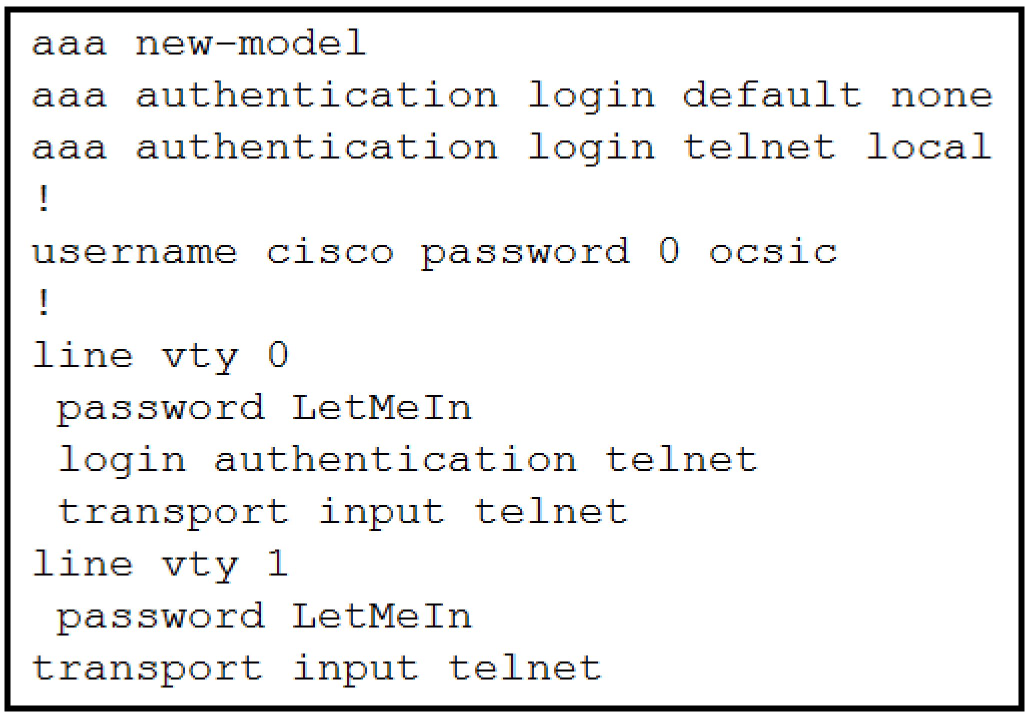

```
aaa new-model
aaa authentication login default none
aaa authentication login telnet local
!
username cisco password 0 ocsic
!
line vty 0
 password LetMeIn
 login authentication telnet
 transport input telnet
line vty 1
 password LetMeIn
transport input telnet
```

Refer to the exhibit. Drag and drop the credentials from the left onto the remote login information on the right to resolve a failed login attempt to vtys. Not all credentials are used.

Select and Place:

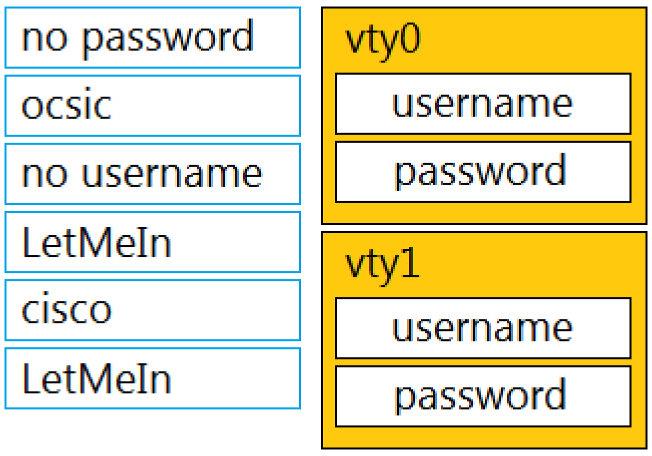

Correct Answer:

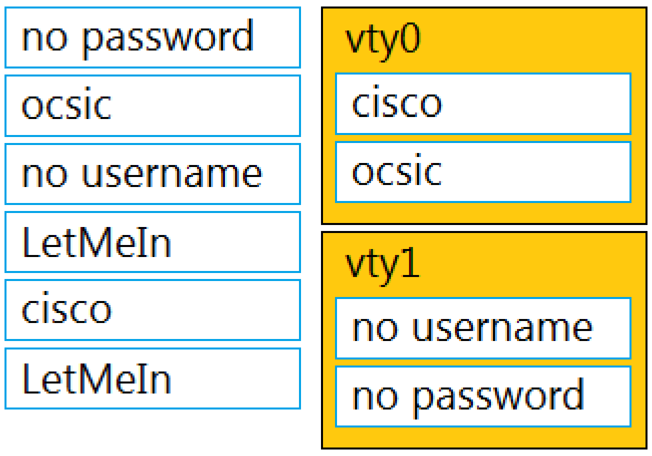

AliMo123 Highly Voted 2 years, 1 month ago

answers are correct
"The command “aaa authentication login default none” means no authentication is required when access to the device via Console/VTY/AUX so if one interface does not specify another login authentication method (via the “login authentication …” command), it will allow to access without requiring username or password. In this case VTY 1 does not specify another authentication login method so it will use the default method (which is “none” in this case)."

upvoted 7 times

Calyfas Most Recent 9 months, 3 weeks ago

Given answer is correct.

upvoted 1 times

Alexloh 11 months ago

When you enable aaa new-model the command 'login authentication default' gets applied to the vty lines. “aaa authentication login default none” also meant no authentication require.

upvoted 2 times

Eric0_0 1 year, 9 months ago

Given answer is correct.
If aaa new-model is NOT configured. The vty password will become effective.

upvoted 2 times

studybuddy10 2 years, 1 month ago

tested on 15.7 code - given answer is correct, user pass on VTY0, no user or pass on VTY1

upvoted 2 times

Baderkhalouf 2 years, 1 month ago

In line VTY 1, there is a password LetMeIn.
Test on Cisco routers.

upvoted 2 times

> **AliMo123** 2 years, 1 month ago
>
> under VTY1, there is no login command, so it goes with default which is None in aaa authentication command
>
> upvoted 5 times

MP_iBGP 2 years, 2 months ago

Answer is correct !

upvoted 1 times

OakA1 2 years, 2 months ago

This is a good link for a reference https://www.cisco.com/c/en/us/support/docs/ios-nx-os-software/authentication-authorization-accounting-aaa/200173-Verify-AAA-behaviour-when-login-local.pdf

upvoted 1 times

Question #153 Topic 1

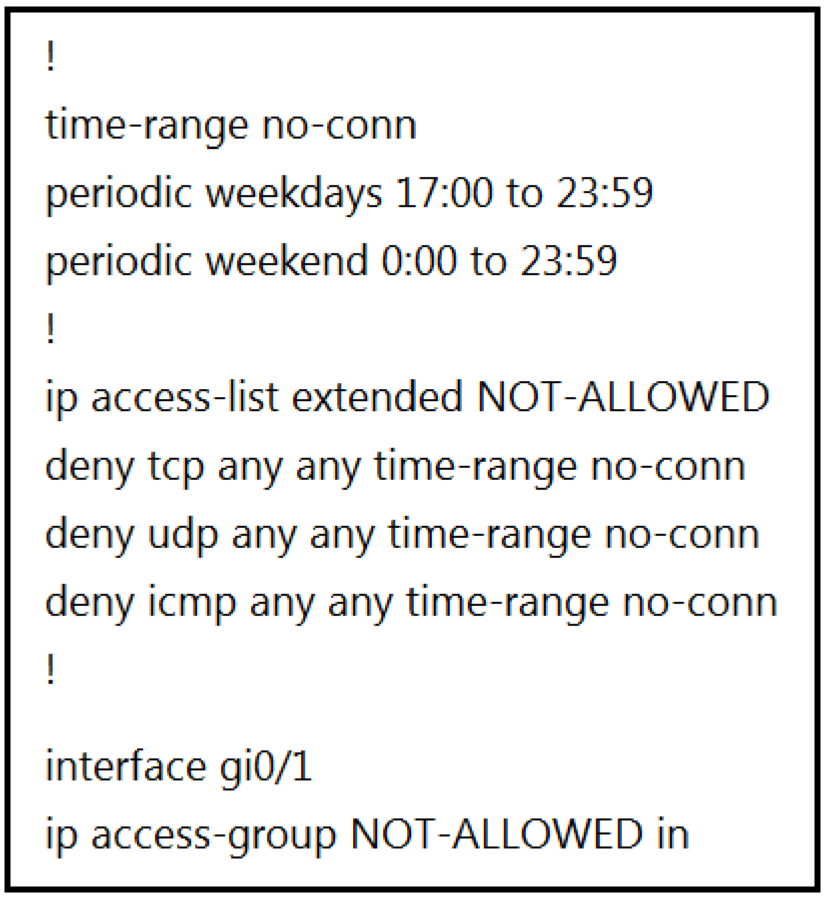

```
!
time-range no-conn
periodic weekdays 17:00 to 23:59
periodic weekend 0:00 to 23:59
!
ip access-list extended NOT-ALLOWED
deny tcp any any time-range no-conn
deny udp any any time-range no-conn
deny icmp any any time-range no-conn
!

interface gi0/1
ip access-group NOT-ALLOWED in
```

Refer to the exhibit. A network administrator wants to block all traffic toward the Internet after business hours and on weekends. When the administrator applies an access list on interface Gi0/1, all traffic is blocked and there is no access to the Internet at any time.
Which action resolves the issue?

A. Add the permit ip any any time-range no-conn statement after the deny udp any any time-range no-conn command in the access list.

B. Add the permit ip any any statement after the deny icmp any any time-range no-conn command in the access list.

C. Add the permit allowed time-range no-conn statement after the deny icmp any any time-range no-conn command in the access list.

D. Add the permit ip any any time-range no-conn statement after the deny icmp any any time-range no-conn command in the access list.

Correct Answer: *B*

Community vote distribution

B (100%)

leecharxos Highly Voted 1 year, 11 months ago

Selected Answer: B

without the statement "permit ip any any" wins the default line of every ACL deny all

upvoted 6 times

examShark Most Recent 2 years, 4 months ago

The given answer is correct

upvoted 3 times

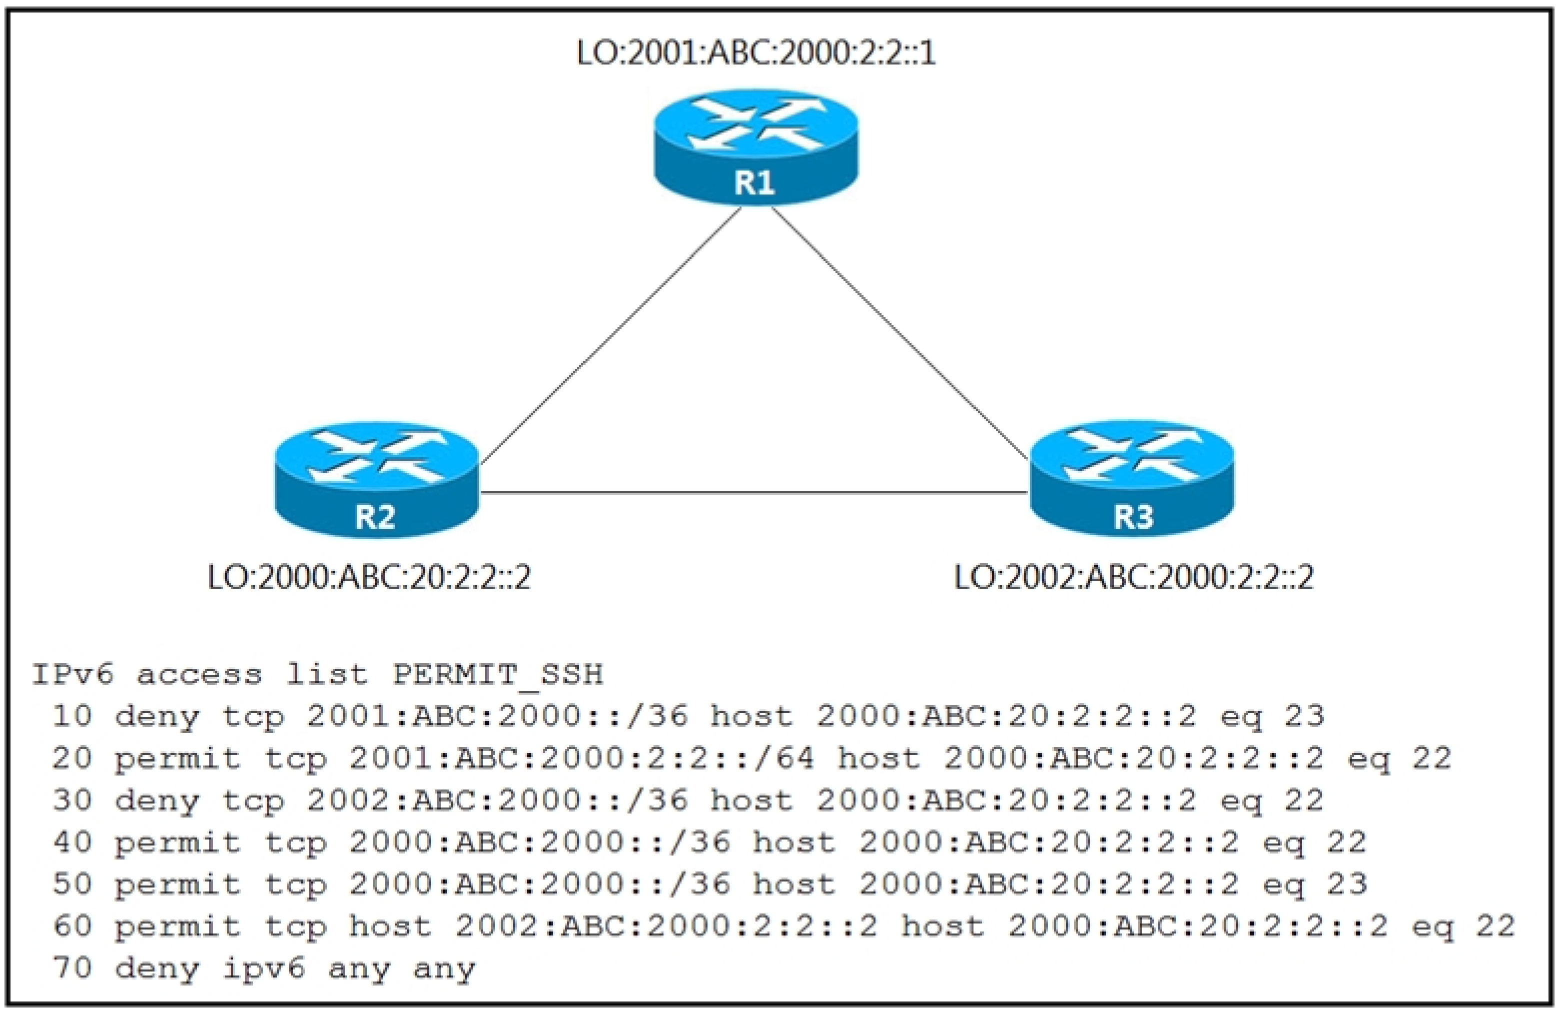

```
IPv6 access list PERMIT_SSH
 10 deny tcp 2001:ABC:2000::/36 host 2000:ABC:20:2:2::2 eq 23
 20 permit tcp 2001:ABC:2000:2:2::/64 host 2000:ABC:20:2:2::2 eq 22
 30 deny tcp 2002:ABC:2000::/36 host 2000:ABC:20:2:2::2 eq 22
 40 permit tcp 2000:ABC:2000::/36 host 2000:ABC:20:2:2::2 eq 22
 50 permit tcp 2000:ABC:2000::/36 host 2000:ABC:20:2:2::2 eq 23
 60 permit tcp host 2002:ABC:2000:2:2::2 host 2000:ABC:20:2:2::2 eq 22
 70 deny ipv6 any any
```

Refer to the exhibit. An IPv6 network was newly deployed in the environment, and the help desk reports that R3 cannot SSH to the R2s Loopback interface.

Which action resolves the issue?

A. Modify line 10 of the access list to permit instead of deny.

B. Remove line 60 from the access list.

C. Modify line 30 of the access list to permit instead of deny.

D. Remove line 70 from the access list.

Correct Answer: *C*

Community vote distribution

C (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: C

correct C

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: C

network range:
2002:0abc:2000:0000:0000:0000:0000:0000 - 2002:0abc:2fff:ffff:ffff:ffff:ffff:ffff

upvoted 1 times

Calyfas 9 months, 3 weeks ago

Given answer is correct.

upvoted 1 times

GreatDane 1 year, 4 months ago

Line 30 of the ACL denies SSH (port 22) traffic from subnet 2002:ABC:2000::/36 to host 2000:ABC:20:2:2::2 (R2).
Here is the problem.

A. Modify line 10 of the access list to permit instead of deny.

Wrong answer.

B. Remove line 60 from the access list.

Wrong answer.

C. Modify line 30 of the access list to permit instead of deny.

Correct answer.

D. Remove line 70 from the access list.

Wrong answer.

upvoted 1 times

Bigmikemalta 1 year, 8 months ago

Selected Answer: C

Given answer is correct

upvoted 1 times

enterTheDevOps 1 year, 11 months ago

I haven't done a ton of IPv6, but... how is the answer correct? It looks to be permitted one way, but return traffic is denied. shouldn't the modification have "permit tcp R2 eq 22 r3"?

upvoted 1 times

enterTheDevOps 1 year, 11 months ago

Cancel that, I see it. I was mistaken. The given answer is correct

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

Question #155

Topic 1

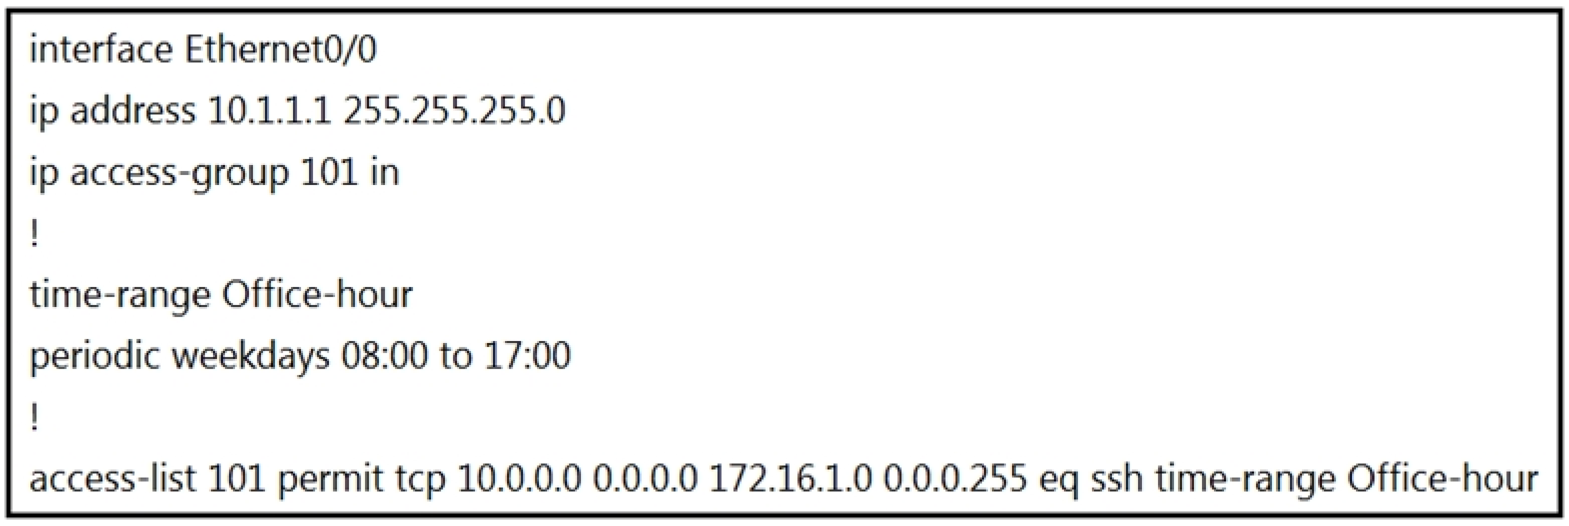

```
interface Ethernet0/0
ip address 10.1.1.1 255.255.255.0
ip access-group 101 in
!
time-range Office-hour
periodic weekdays 08:00 to 17:00
!
access-list 101 permit tcp 10.0.0.0 0.0.0.0 172.16.1.0 0.0.0.255 eq ssh time-range Office-hour
```

Refer to the exhibit. An IT staff member comes into the office during normal office hours and cannot access devices through SSH.
Which action should be taken to resolve this issue?

A. Modify the access list to use the correct IP address.

B. Configure the correct time range.

C. Modify the access list to correct the subnet mask.

D. Configure the access list in the outbound direction.

Correct Answer: *C*

Community vote distribution

A (63%) C (37%)

cakmamail Highly Voted 2 years, 4 months ago

I changed my mind, i think it is A.
Because C says subnetmask. And i dont think they would use the word subnet mask instead of wildcard mask.
For A to be true, we need to know that IT guy`s ip address and use that to correct the ACL

upvoted 10 times

BTK0311 Most Recent 3 months ago

Selected Answer: C

permit 10.0.0.0 0.0.0.0 will only allow a host with 10.0.0.0 IP but subnet is the wrong word, should be mask.

upvoted 1 times

jansan55 4 months ago

Selected Answer: C

My choice: Answer C
Enough to change the ACL like this:
access-list 101 permit tcp 10.0.0.0 0.1.255.255 172.16.1.0 0.0.0.255 eq ssh time-range Office-hour
With answer A, we get only one IP address, from where ssh allowed, while this company has an IT staff.

upvoted 1 times

[Removed] 4 months, 1 week ago

Selected Answer: C

Okay, I will go with C. I was torn between A and C, but C seems more plausible as the answer because chainging the IP address of the source portion of the ACL will only apply to one host device, when there could be a Staff with multiple devices...
I agree that there may be a discrepancy in wording of Subnet Mask and Wildcard mask, but subnetmask can be changed from 0.0.0.0 to 0.255.255.255 to cover the correct subnetmask.

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: A

The source 10.0.0.0 0.0.0.0 means host 10.0.0.0, and it is not valid for this topology.
So, we need to correct the source ip address for sure.

upvoted 3 times

HungarianDish 6 months, 3 weeks ago

The information is missing, what should we set as the source in the ACL.
Is the device shown in the question the source or the destination of the telnet traffic? Or is telnet transiting through it?

If it is the source, and telnet should be initiated from this device (10.1.1.1 0.0.0.0) to other devices (172.16.1.0 0.0.0.255), then:
-the ACL won't work. We can't apply any ACL to the outbound traffic generated locally by the router itself

If telnet is transiting through this device (for instance, coming from a LAN connected to E0/0), then:
-we should correct the ip address and wildcard mask, too:
access-list 101 permit tcp 10.1.1.0 0.0.0.255 172.16.1.0 0.0.0.255 eq ssh time-range Office-hour

The device with IP 10.1.1.1 could also be the destination, and telnet traffic would enter on E0/0 inbound. In that case the ACL would be something like this:
access-list 101 permit tcp 172.16.1.0 0.0.0.255 host 10.1.1.1 eq ssh time-range Office-hour

The output does not show clearly, how they want to use the ACL.
upvoted 2 times

HungarianDish 6 months, 3 weeks ago

*I meant SSH traffic.
upvoted 1 times

Malasxd 7 months, 1 week ago

I would chose "C", but the word "subnet mask" got me...
"A" seems more right, but I am not sure.
upvoted 2 times

Dacusai 7 months, 3 weeks ago

A
A is more accurate but you have to modify both IP and Wilcard 10.1.1.0 0.0.0.255 it should be like that
upvoted 1 times

Alexloh 11 months ago

Selected Answer: A

I believed (A) is correct answer, below is the intended config:

access-list 101 permit tcp 10.1.1.1 0.0.0.0 172.16.1.0 0.0.0.255 eq ssh time-range Office-hour
upvoted 3 times

CisconAWSGURU 1 year, 1 month ago

Selected Answer: A

I like A, more
upvoted 2 times

Huntkey 1 year, 2 months ago

Selected Answer: C

The question didn't say what IP the connection is from or to. It didn't say the SSH is to the router itself. It is more than likely the SSH traffic is through the router instead of destined or sourced from the router. In that case, I think C would make more sense. 10.0.0.0/0.0.0.0 is clearly wrong.
upvoted 1 times

GreatDane 1 year, 4 months ago

On a router, access-list 101 permits SSH connections from 10.0.0.0/0.0.0.0, which equals to 10.0.0.0/255.255.255.255, which equals to 10.0.0.0/32. In other words, SSH access is allowed only to this IP address, and not to a subnet.
The correct syntax could be:

access-list 101 permit tcp 10.0.0.0 0.0.0.255 ...

But, since the question refers to a single IT staff member, the solution to the problem could be allowing SSH access only to a single IP address, like this:

access-list 101 permit tcp 10.1.1.1 0.0.0.0 ...

A. Modify the access list to use the correct IP address.

Correct answer.

B. Configure the correct time range.

Wrong answer.

C. Modify the access list to correct the subnet mask.

Wrong answer.

D. Configure the access list in the outbound direction.

Wrong answer.
upvoted 1 times

TECH3K3 1 year, 5 months ago

This is such a BS question as you could change the IP address or subnet mask, which is really the wildcard. I be selecting A to change the IP and leave the wildcard to match all bits
upvoted 1 times

kellyDD 1 year, 6 months ago

It should be the wildcard mask, not the subnet mask, that is set in the access list.

upvoted 1 times

kellyDD 1 year, 6 months ago

The problem before that is that the access list settings are wrong to begin with, not just the wildcard, so that needs to be corrected. Therefore, the answer is A.

upvoted 1 times

kellyDD 1 year, 4 months ago

I hate this question.Oh no - I don't want to do it, I don't want to do it. This output screen and the choices are a mess. wildcard is a wildcard and subnet mask is a subnet mask.

upvoted 1 times

cyrus777 1 year, 8 months ago

we don't have subnet mask in here. A is correct

upvoted 3 times

wts 1 year, 9 months ago

Selected Answer: A

..the main problem is sender address 10.0.0.0 It's unlikely that our worker has such an address configured. And then we should choose honey A and C.

A - if it is assumed that the employee works from a PC from the network 10.1.1.0/24, then changing the address to 10.1.1.x/32 is reasonable.

C - let's say we set /8. It's not very elegant, but any package from 10.1.1.0/24 will pass this access list.

It seems to me that opting for a stricter rule is more correct than giving access to the entire 10/8 network.

upvoted 4 times

Networkingguy 1 year, 10 months ago

Real 50/50 split on this one, A and C both are correct, but C would be more correct as old mate JingleJangus has pointed out it would be better to have the whole range for all IT staff members.

upvoted 1 times

JingleJangus 1 year, 10 months ago

Selected Answer: C

C is more correct because even if A were true, it would only allow the default gateway to access devices via ssh using the source int eth0/0. Because changing 10.0.0.0 0.0.0.0 to ---> 10.1.1.1 or even 10.1.1.0 0.0.0.0.0 would only allow 10.1.1.1 out, or idek if 10.1.1.0 /32 would work. But yeah C because if you modify the ACL to 10.0.0.0 0.255.255.255, it will still allow the traffic to pass thru via ssh beyond the DG.

upvoted 3 times

Question #156 Topic 1

Refer to the exhibit.

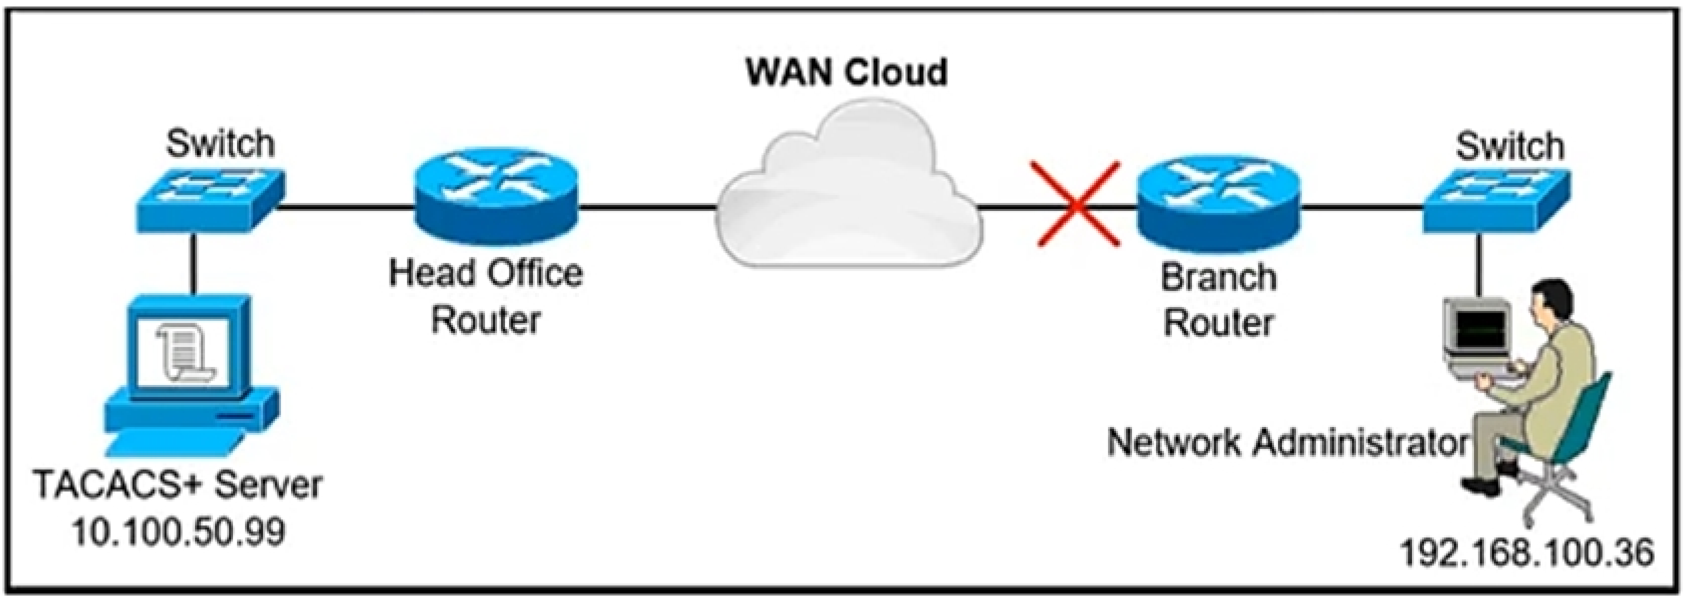

A network administrator is trying to access a branch router using TACACS+ username and password credentials, but the administrator cannot log in to the router because the WAN connectivity is down. The branch router has following AAA configuration: aaa new-model aaa authorization commands 15 default group tacacs+ aaa accounting commands 1 default stop-only group tacacs+ aaa accounting commands 15 default stop-only group tacacs+ tacacs-server host 10.100.50.99 tacacs-server key Ci$co123
Which command will resolve this problem when WAN connectivity is down?

A. aaa authentication login console group tacacs+ enable

B. aaa authentication login default group tacacs+ local

C. aaa authentication login default group tacacs+ enable

D. aaa authentication login default group tacacs+ console

Correct Answer: *B*
Reference:
https://www.cisco.com/c/en/us/support/docs/security-vpn/terminal-access-controller-access-control-system-tacacs-/200606-aaa-authentication-login- default-local.html

Community vote distribution

B (100%)

bogd Highly Voted 1 year, 10 months ago
Selected Answer: B

```
aaa new-model
aaa authorization commands 15 default group tacacs+
aaa accounting commands 1 default stop-only group tacacs+
aaa accounting commands 15 default stop-only group tacacs+
tacacs-server host 10.100.50.99
tacacs-server key Ci$co123
```

Both B and C would work (we do not see the rest of the config, we do not know whether users or enable secrets are configured)
upvoted 7 times

Huntkey Most Recent 1 year, 2 months ago
I just don't know how it would work when the WAN is up without the "aaa authentication login" configuration. Does it by default uses TACACS for authentication?
upvoted 2 times

HungarianDish 6 months, 3 weeks ago
probably, this is already configured "aaa authentication login default group tacacs+"
upvoted 1 times

GreatDane 1 year, 4 months ago
If the TACACS+ server is unavailable, the only way to log on to the router is to enable local authentication.

A. aaa authentication login console group tacacs+ enable

Wrong answer.

B. aaa authentication login default group tacacs+ local

Correct answer.

C. aaa authentication login default group tacacs+ enable

Wrong answer.

D. aaa authentication login default group tacacs+ console

Wrong answer.

upvoted 1 times

Refer to the exhibit.

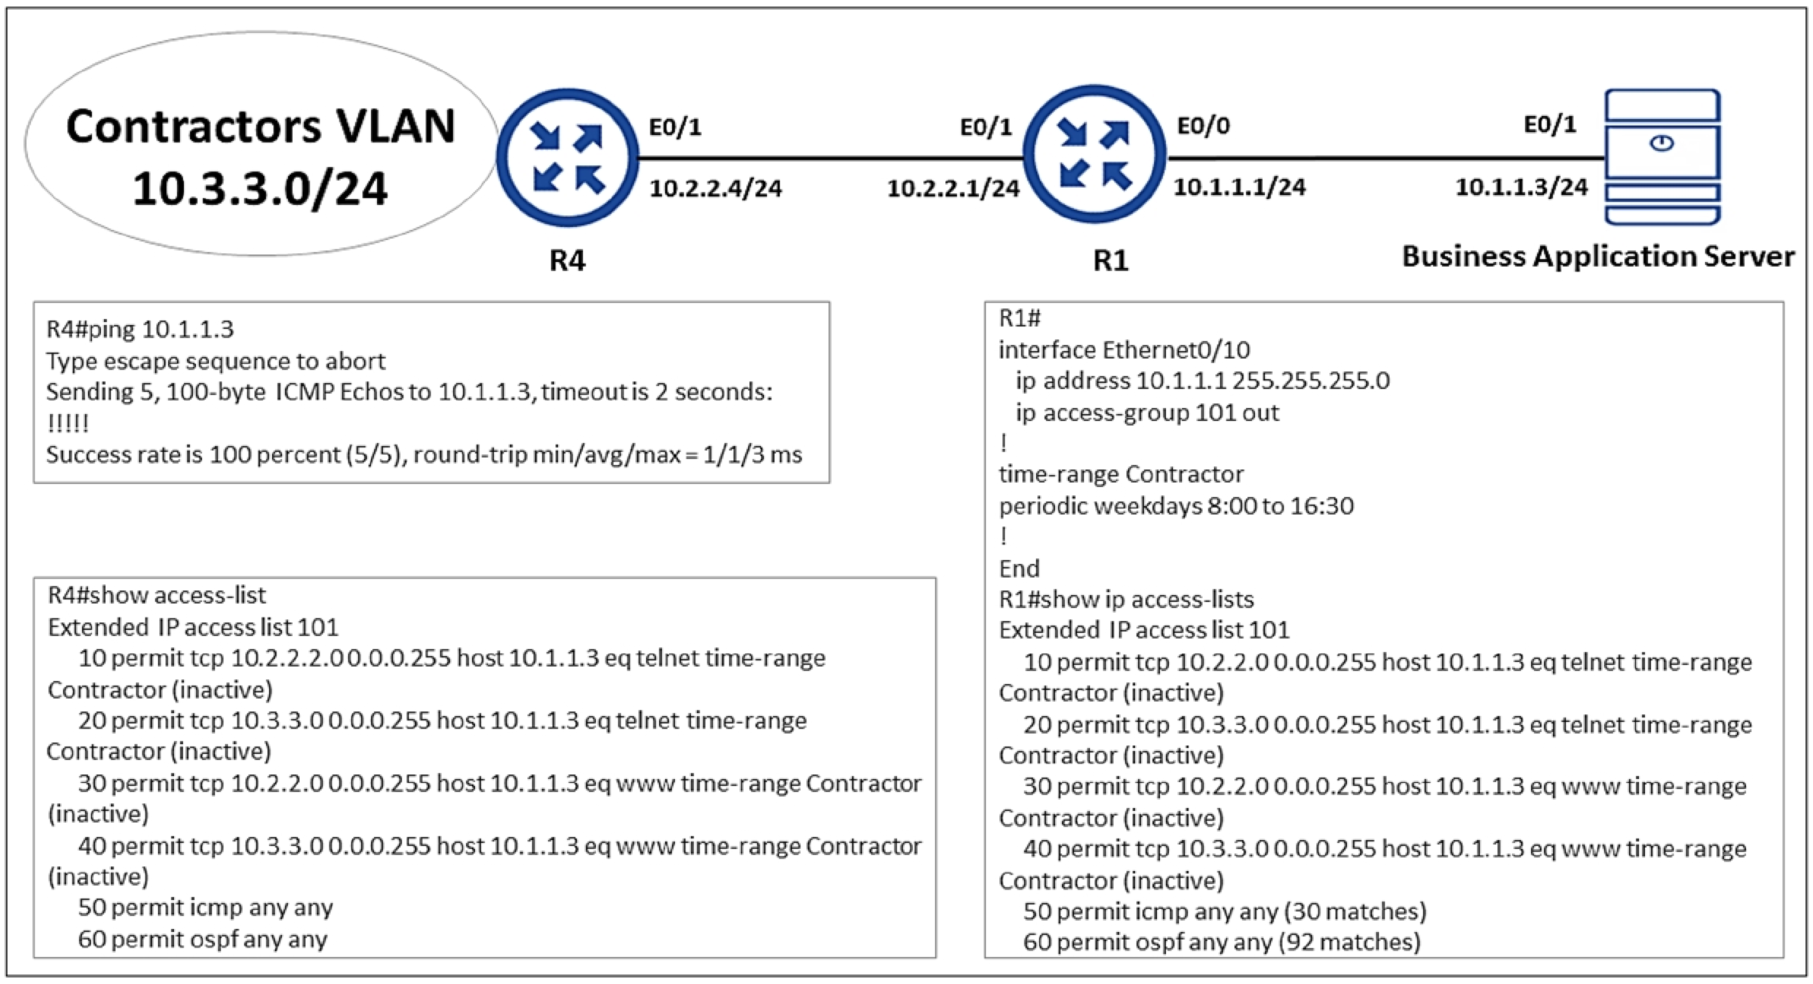

```
R4#ping 10.1.1.3
Type escape sequence to abort
Sending 5, 100-byte ICMP Echos to 10.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/3 ms
```

```
R4#show access-list
Extended IP access list 101
   10 permit tcp 10.2.2.2.0 0.0.0.255 host 10.1.1.3 eq telnet time-range
Contractor (inactive)
   20 permit tcp 10.3.3.0 0.0.0.255 host 10.1.1.3 eq telnet time-range
Contractor (inactive)
   30 permit tcp 10.2.2.0 0.0.0.255 host 10.1.1.3 eq www time-range Contractor
(inactive)
   40 permit tcp 10.3.3.0 0.0.0.255 host 10.1.1.3 eq www time-range Contractor
(inactive)
   50 permit icmp any any
   60 permit ospf any any
```

```
R1#
interface Ethernet0/10
  ip address 10.1.1.1 255.255.255.0
  ip access-group 101 out
!
time-range Contractor
periodic weekdays 8:00 to 16:30
!
End
R1#show ip access-lists
Extended IP access list 101
   10 permit tcp 10.2.2.0 0.0.0.255 host 10.1.1.3 eq telnet time-range
Contractor (inactive)
   20 permit tcp 10.3.3.0 0.0.0.255 host 10.1.1.3 eq telnet time-range
Contractor (inactive)
   30 permit tcp 10.2.2.0 0.0.0.255 host 10.1.1.3 eq www time-range
Contractor (inactive)
   40 permit tcp 10.3.3.0 0.0.0.255 host 10.1.1.3 eq www time-range
Contractor (inactive)
   50 permit icmp any any (30 matches)
   60 permit ospf any any (92 matches)
```

An engineer is troubleshooting failed access by contractors to the business application server via Telnet or HTTP during the weekend.
Which configuration resolves the issue?

A. R1 no access-list 101 permit tcp 10.3.3.0 0.0.0.255 host 10.1.1.3 eq telnet time-range Contractor

B. R1 time-range Contractor no periodic weekdays 8:00 to 16:30 periodic daily 8:00 to 16:30

C. R4 time-range Contractor no periodic weekdays 17:00 to 23:59 periodic daily 8:00 to 16:30

D. R4 no access-list 101 permit tcp 10.3.3.0 0.0.0.255 host 10.1.1.3 eq telnet time-range Contractor

Correct Answer: *B*

Community vote distribution

B (88%) 13%

idlechado 2 months, 3 weeks ago

C is wrong:

ACL 101 in R1 has not matches which means ACL 101 in R1 is not inside the required flow. Remember, they show a successful ping, and ACL 101 in R4 has matches of this ping. Therefor it's not necessary to change anything in R1

upvoted 2 times

[Removed] 4 months, 1 week ago

B is correct,
A and D are irrelevant, the ACEs are correct in regards to source and destination.
C is wrong because we don't know whether R4 has the correct schedule, we can only assume it does because the exhibit only displays R1's schedule and that is not covering Weekend Days, only weekdays, therefore we have to remove it, and include saturday and sunday.

upvoted 2 times

inteldarvid 5 months ago

Selected Answer: B

yes, option B correct

https://www.cisco.com/E-Learning/bulk/public/tac/cim/cib/using_cisco_ios_software/cmdrefs/periodic.htm

upvoted 2 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: B

see explanation from daloslav

upvoted 2 times

daloslav 7 months ago

Selected Answer: B

Contractors need to connect during weekend but time-range is configured for weekdays. Answer B is correct because you have to delete old time-range statement for weekdays, and configure new for all days (daily).
C is not correct because bad time-range (17:00 to 23:59).

upvoted 3 times

Malasxd 7 months, 1 week ago

Selected Answer: C

It's B or C.

The ACL 101 in R1 is applied in outbound direction in interface e0/10 and the interface connected to the server ins e0/0 BUTTTT, in the commands it show the same IP address applied in interface e0/10 that is showed in topology. It's strange.

upvoted 1 times

Malasxd 7 months, 1 week ago

Forget what i said. C does not make any sense. The topology and the command interface in R4 are different (e0/0 and e0/10). I think it is a typing error. B is the only one make sense.

upvoted 2 times

GreatDane 1 year, 4 months ago

This happens because the keyword weekdays in time-range Contractors, on R1, means “Monday through Friday”. To include weekend days, use the keyword daily.

A. R1 no access-list 101 permit tcp 10.3.3.0 0.0.0.255 host 10.1.1.3 eq telnet time-range Contractor

Wrong answer.

B. R1 time-range Contractor no periodic weekdays 8:00 to 16:30 periodic daily 8:00 to 16:30

Correct answer.

C. R4 time-range Contractor no periodic weekdays 17:00 to 23:59 periodic daily 8:00 to 16:30

Wrong answer.

D. R4 no access-list 101 permit tcp 10.3.3.0 0.0.0.255 host 10.1.1.3 eq telnet time-range Contractor

Wrong answer.

upvoted 4 times

wts 1 year, 8 months ago

Contractors cannot get to the server on weekends.
We have extended the time-range (in which this situation occurs) to every day.
...well, OK.

upvoted 2 times

Question #158 Topic 1

What are two characteristics of IPv6 Source Guard? (Choose two.)

A. requires the user to configure a static binding

B. used in service provider deployments to protect DDoS attacks

C. requires that validate prefix be enabled

D. requires IPv6 snooping on Layer 2 access or trunk ports

E. recovers missing binding table entries

Correct Answer: *AD*

Community vote distribution

AD (58%) CE (21%) 8% 8%

HungarianDish Highly Voted 7 months, 3 weeks ago

Selected Answer: AD

This is how I see it: For source guard to operate, binding table entries need to exists. So, A or D are required.
A) static binding -> yes, or use ipv6 snooping #security-level glean to populate the binding table
B) to protect against DDOS -> yes, but not just for service providers (it's rather prefix guard)
C) can be configured with validate address or validate prefix (not explicitly needed)
D) snooping on L2 access or trunk -> yes, or create static bindings
E) not source guard itself, but the snooping feature glean recovers missing binding table entries

upvoted 10 times

Tedmus Most Recent 3 weeks, 6 days ago

Selected Answer: BD

From ENARSI course:
B | Protect against DoS attacks - not only with Service Providers but of course they can use it.
D | IPv6 Snooping is a prerequisite for IPv6 to work.

Not A: The user REQUIRES is wrong. It is possible fo the admin to configure a static binding. But usually it is learned with DHCPv6 or ND.

upvoted 2 times

Pietjeplukgeluk 3 weeks, 3 days ago

I actually agree here the "requires" is wrong. Anyway, i think if you look at this question, the "requires" in answer D is also wrong. A better way of saying: "needs a binding table entry, that could be statically configured", "needs a binding table entry, that can by dynamically configured using snooping on L2 access or trunk". Concluding, i still think A and D is best, B could be accurate, but i don't work for any provider, they could rely on different technologies also to filter inbound traffic on correct source.

upvoted 1 times

chris110 3 months, 1 week ago

Selected Answer: AC

IPv6 Source Guard uses the IPv6 First-Hop Security Binding Table to drop traffic from unknown sources or bogus IPv6 addresses not in the binding table. The switch also tries to recover from lost address information, querying DHCPv6 server or using IPv6 neighbor discovery to verify the source IPv6 address after dropping the offending packet(s).

Reference: https://blog.ipspace.net/2013/07/first-hop-ipv6-security-features-in.html

Although IPv6 Source Guard looks at information in the binding table and IPv6 snooping can fill this table but IPv6 snooping is not a must to run IPv6 Source Guard. We can use other methods to fill the binding table like static binding or ND inspection -> Answer 'requires IPv6 snooping on Layer 2 access or trunk ports' is not correct.

IPv6 Source Guard is used to mitigate attacks from hosts connected to untrusted access interfaces on the switch -> Answer 'used in service provider deployments to protect DDoS attacks' is not correct.

Answer 'requires the user to configure a static binding' is not correct as we can use IPv6 Snooping feature to populate the IPv6 binding table.

upvoted 1 times

chris110 3 months, 1 week ago

i mean c & e

upvoted 1 times

gpaulino 4 months, 2 weeks ago

Selected Answer: AD

IPv6 Source Guard is a feature that enhances network security by ensuring that the source IPv6 addresses in incoming packets are valid and legitimate. It helps prevent spoofing attacks and unauthorized address usage. Among the options you've provided, the following are the two correct characteristics of IPv6 Source Guard:

A. Requires the user to configure a static binding.

This is correct. IPv6 Source Guard can work in conjunction with IPv6 snooping to create a binding table of legitimate IPv6 addresses associated with specific Layer 2 ports. The administrator can manually configure static bindings to explicitly define which IPv6 addresses are allowed to originate from specific ports.
D. Requires IPv6 snooping on Layer 2 access or trunk ports.

This is correct. IPv6 Source Guard relies on IPv6 snooping to build and maintain a binding table that correlates IPv6 addresses with their corresponding Layer 2 ports. By snooping on Layer 2 traffic, the switch can learn and enforce valid bindings between IPv6 addresses and physical interfaces.
The other options (B, C, and E) are not accurate characteristics of IPv6 Source Guard
upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: AD

A and D
upvoted 1 times

OskarNorman 6 months, 4 weeks ago

It is C and E
upvoted 1 times

MasterMatt 8 months, 3 weeks ago

Selected Answer: CE

Answer is CE
upvoted 1 times

Zizu007 11 months, 2 weeks ago

Selected Answer: AD

Answer is Correct!
IPv6 Source Guard is a "Data-plane" filter --> creates automatically IPv6 PACL to filter sources.

This automatic PACL is used ingress on a port. And it uses one or more sources;
- IPv6 snooping;
- DHCPv6 or NDP RA/RS msgs
- Static entries

Static entry is required for the attached device who has static IPv6 addresses configured (router/printer/server)
upvoted 2 times

PimplePooper 12 months ago

Selected Answer: CE

Answer is CE
upvoted 2 times

Ckl22 1 year ago

Selected Answer: CD

IPv6 source guard is an interface between the populated binding table and data traffic filtering, and the binding table must be populated with IPv6 prefixes for IPv6 source guard to work.

IPv6 Source Guard and IPv6 Prefix Guard are Layer 2 snooping features that validate the source of IPv6 traffic

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-3s/ip6f-xe-3s-book/ip6-src-guard.html
upvoted 1 times

GreatDane 1 year, 4 months ago

A. requires the user to configure a static binding

IPv6 Source Guard relies on DHCP and ND protocols. A static binding can be configured in the snooping table, but it’s not required.
Wrong answer.

B. used in service provider deployments to protect DDoS attacks

Something like Cisco Guard XT.
Wrong answer.

C. requires that validate prefix be enabled

This is IPv6 Prefix Guard configuration: enables IPv6 Source Guard to perform the IPv6 Prefix-Guard operation.
Correct answer.

D. requires IPv6 snooping on Layer 2 access or trunk ports

Wrong answer.

E. recovers missing binding table entries

This is the IPv6 First-Hop Security Binding Table Recovery Mechanism.
Correct answer.

upvoted 2 times

cisconut 1 year, 5 months ago

Selected Answer: CE

Cisco doc says "When traffic is denied, the IPv6 address glean feature is notified so
that it can try to recover the traffic by querying the DHCP server or by using IPv6 ND.".

upvoted 1 times

timtgh 1 year, 6 months ago

Selected Answer: CD

Confirmed in Cisco docs.

upvoted 1 times

xziomal9 1 year, 7 months ago

Selected Answer: CE

The correct answer is: C E

upvoted 1 times

JOKERR 1 year, 6 months ago

It's not E.

Source Guard only looks at information found in the binding table, and it doesn't fill the binding table. You need another feature like ND inspection or IPv6 snooping to do this.

upvoted 4 times

Question #159 *Topic 1*

DRAG DROP -

Drag and drop the IPv6 first hop security device roles from the left onto the corresponding descriptions on the right.

Select and Place:

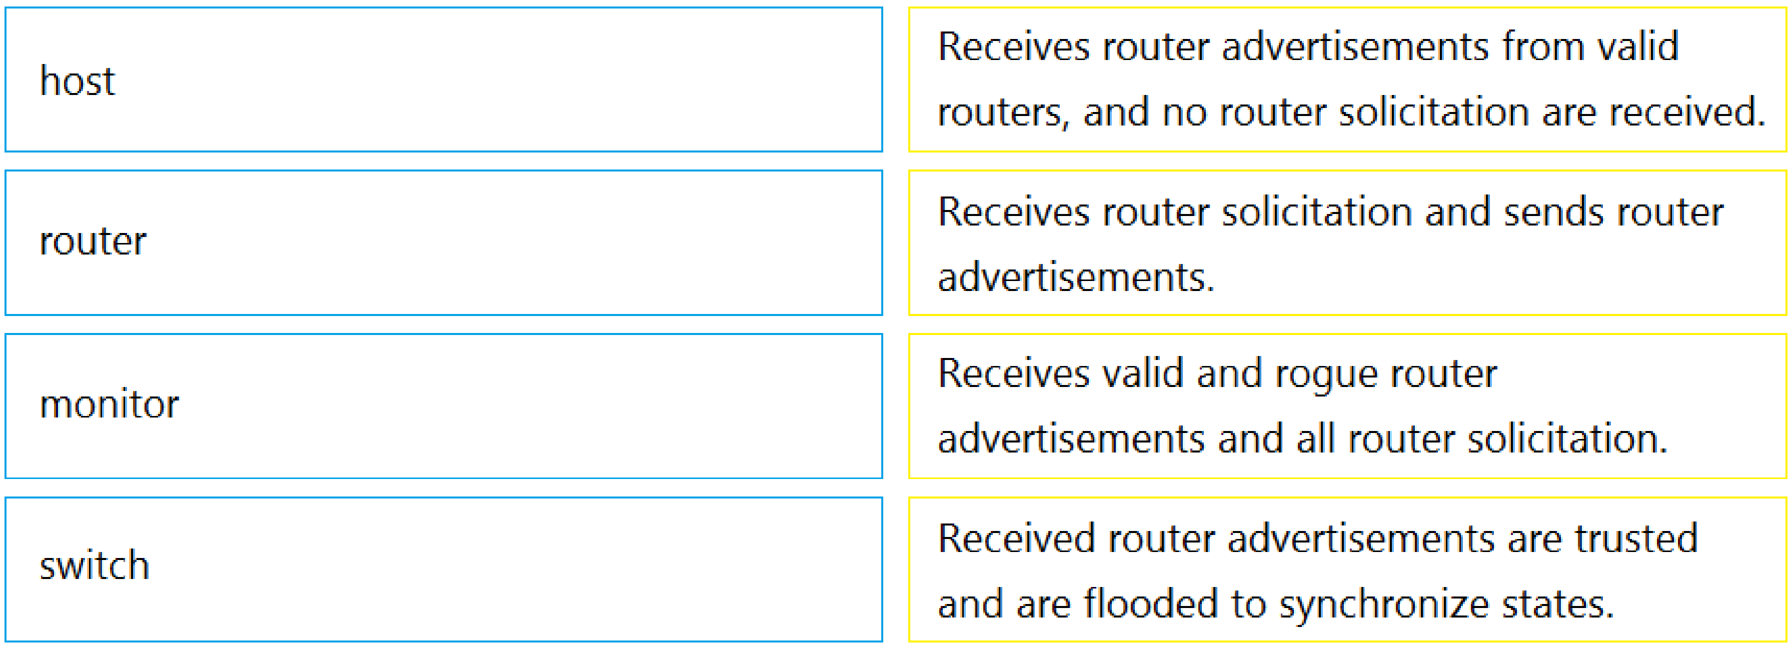

host	Receives router advertisements from valid routers, and no router solicitation are received.
router	Receives router solicitation and sends router advertisements.
monitor	Receives valid and rogue router advertisements and all router solicitation.
switch	Received router advertisements are trusted and are flooded to synchronize states.

Correct Answer:

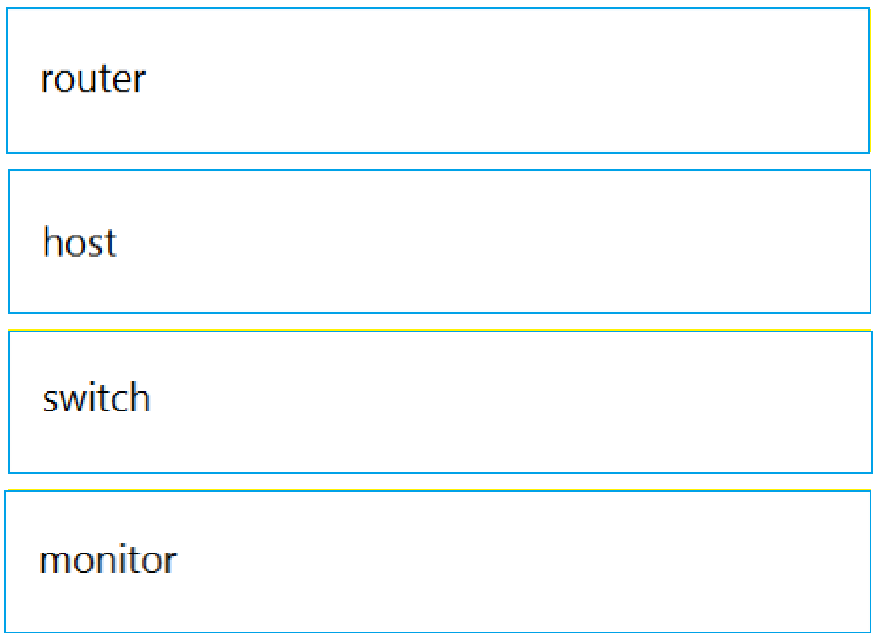

router

host

switch

monitor

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/security/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x_chapter_011011.pdf

t1s Highly Voted 1 year, 5 months ago

Device Roles for RA-guard, devices can have different roles:
• Host (default): can only receive RA from valid routers, no RS will be received
• Router: can receive RS and send RA
• Monitor: receive valid and rogue RA and all RS
• Switch: RA are trusted and flooded to synchronize states

Source:
https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/BRKSEC-3200.pdf

upvoted 12 times

Huntkey Highly Voted 1 year, 2 months ago

so according to below, the answer is incorrect. IT should be
Host
Router
Monitor
Switch

upvoted 6 times

Tedmus Most Recent 3 weeks, 6 days ago

This link explained it:
https://www.ciscolive.com/c/dam/r/ciscolive/global-event/docs/2022/pdf/BRKENT-3002.pdf

Important to consider is the direction. In this case we are talking about the devices itself and not about the configuration of the Switch-Port.

1: host

2: router
3: monitor
4: switch

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

For RA-guard, devices can have different roles
• Host (default): can only receive RA from valid routers, no RS will be received
• Router: can receive RS and send RA
• Monitor: receive valid and rogue RA and all RS
• Switch: RA are trusted and flooded to synchronize states

upvoted 2 times

WAKIDI 1 year, 5 months ago

Please anyone give the link of reference to "this kind of Monitor"

upvoted 1 times

ytsionis 1 year, 6 months ago

I Think tha is the right order

Host Receives router advertisements from valid routers and no router solicitation are received
Router- Receives router solicitation and sends router advertisements
Switch Receives valid and rogue router advertisements and all router solicitation
Monitor Received router advertisements are trusted and are flooded to synchronize states

upvoted 1 times

Question #160 *Topic 1*

The network administrator configured R1 for Control Plane Policing so that the inbound Telnet traffic is policed to 100 kbps. This policy must not apply to traffic coming in from 10.1.1.1/32 and 172.16.1.1/32. The administrator has configured this: access-list 101 permit tcp host 10.1.1.1 any eq 23 access-list 101 permit tcp host 172.16.1.1 any eq 23

```
!
class-map CoPP-TELNET
match access-group 101
!
policy-map PM-CoPP
class CoPP-TELNET
police 100000 conform transmit exceed drop
!
control-plane
service-policy input PM-CoPP
```

The network administrator is not getting the desired results.

Which set of configurations resolves this issue?

A. no access-list 101 access-list 101 deny tcp host 10.1.1.1 any eq 23 access-list 101 deny tcp host 172.16.1.1 any eq 23 access-list 101 permit ip any any

B. control-plane no service-policy input PM-CoPP ! interface Ethernet 0/0 service-policy input PM-CoPP

C. no access-list 101 access-list 101 deny tcp host 10.1.1.1 any eq 23 access-list 101 deny tcp host 172.16.1.1 any eq 23 access-list 101 permit ip any any ! Interface E 0/0 service-policy input PM-CoPP

D. control-plane no service-policy input PM-CoPP service-policy input PM-CoPP

Correct Answer: *A*

Community vote distribution

A (100%)

mgiuseppe86 2 months, 3 weeks ago

Selected Answer: A

a better fitting answer would be
“access-list 101 permit tcp any any eq 23” in order to police all telnet traffic, which is what the questions asks. Otherwise, all traffic is being policed here.

upvoted 1 times

guy276465281819372 3 months, 4 weeks ago

permit ip any any eq 23 would be nice to have

upvoted 1 times

David98898998 6 months, 2 weeks ago

This is a stupid question because the "permit ip any any" is going to police all traffic except for two particular hosts Telnet traffic. It will not do as desired. Still, A is best answer.

upvoted 2 times

Xerath 10 months ago

Selected Answer: A

The given answer is correct.

upvoted 2 times

Ghadir2023 10 months, 3 weeks ago

packets that match a deny rule are excluded from that class and cascade to the next class (if one exists) for classification. Therefore, if we don’t want to CoPP traffic from 10.1.1.1/32 and 172.16.1.1/32, we must “deny” them in the ACL.

upvoted 2 times

GreatDane 1 year, 4 months ago

What’s missing here is the definition of ACL 101.

A. no access-list 101 access-list 101 deny tcp host 10.1.1.1 any eq 23 access-list 101 deny tcp host 172.16.1.1 any eq 23 access-list 101 permit ip any any

This syntax
flushes any previous ACL 101 statement

denies any Telnet traffic from 10.1.1.1/32 and 172.16.1.1/32
permits any other IP traffic

Correct answer.

B. control-plane no service-policy input PM-CoPP ! interface Ethernet 0/0 service-policy input PM-CoPP

Wrong answer.

C. no access-list 101 access-list 101 deny tcp host 10.1.1.1 any eq 23 access-list 101 deny tcp host 172.16.1.1 any eq 23 access-list 101 permit ip any any ! Interface E 0/0 service-policy input PM-CoPP

CoPP policy PM-COPP is already assigned to the control plane context.
Wrong answer.

D. control-plane no service-policy input PM-CoPP service-policy input PM-CoPP

Wrong answer.

upvoted 1 times

WAKIDI 1 year, 5 months ago

A is the correct answer. reference : https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/15-0SY/configuration/guide/15_0_sy_swcg/control_plane_policing_copp.pdf Page 8. This example shows how to allow full access for Telnet to the switch from a host in a specific subnet and police the rest of the subnet:
Router(config)# access-list 121 deny tcp host 10.86.183.3 any eq telnet
Router(config)# access-list 121 permit tcp 10.86.183.0 0.0.0.255 any eq telnet

upvoted 2 times

Question #161 *Topic 1*

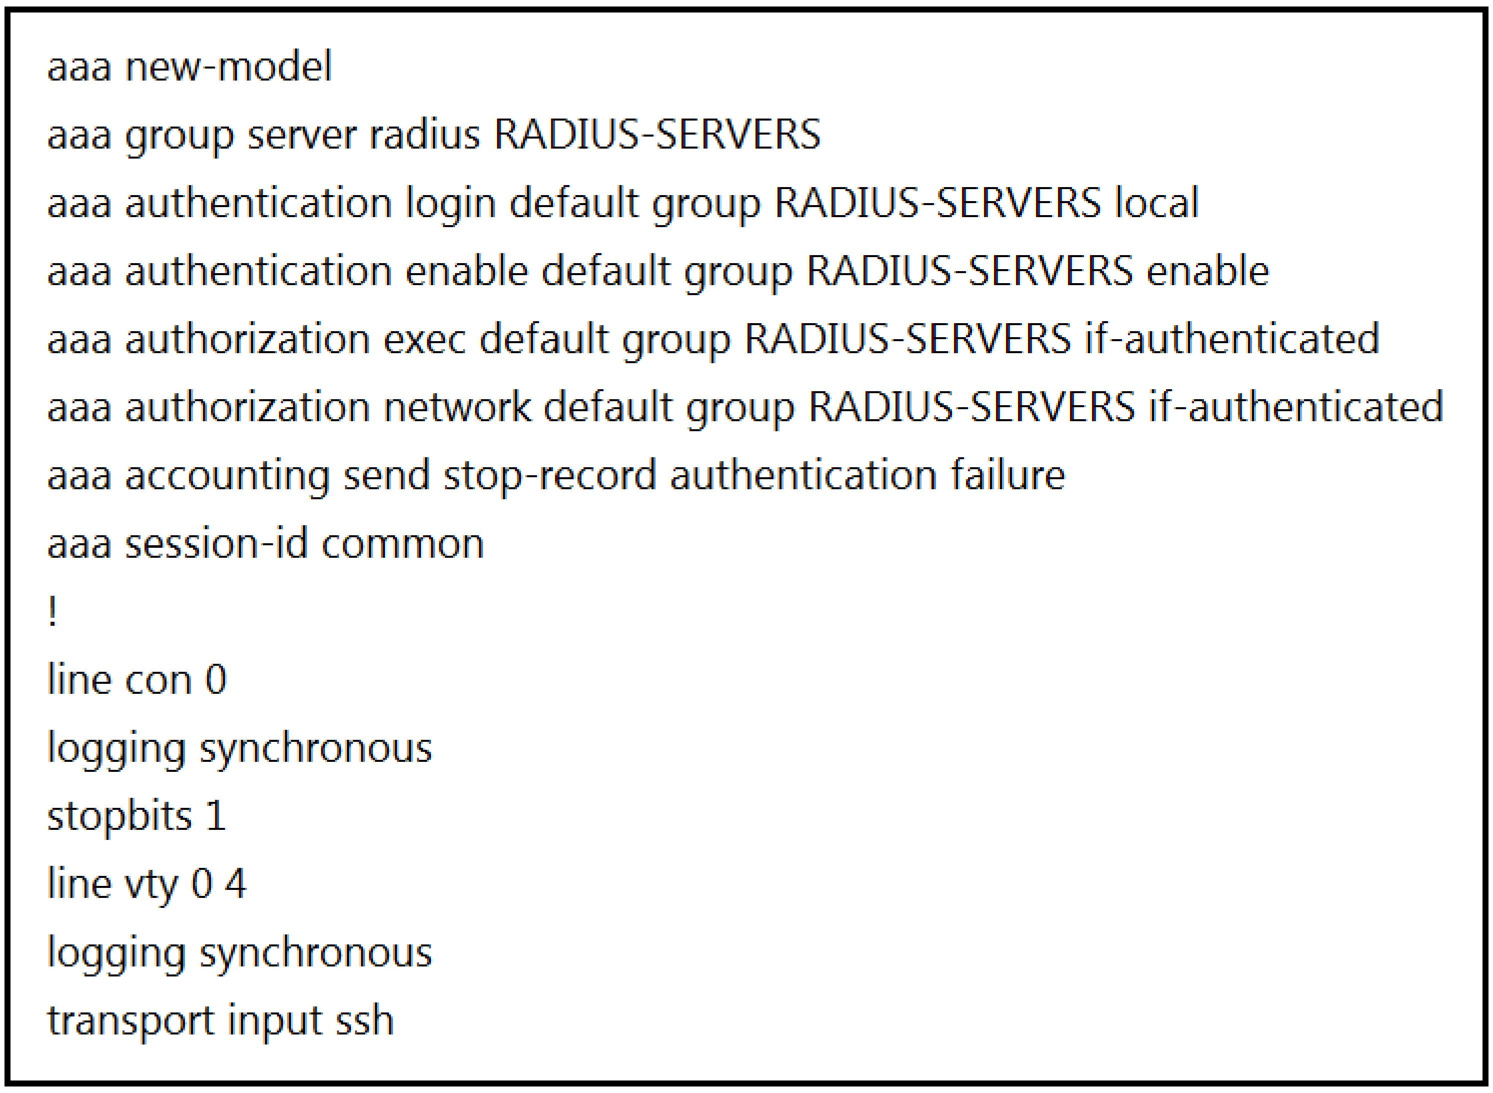

```
aaa new-model
aaa group server radius RADIUS-SERVERS
aaa authentication login default group RADIUS-SERVERS local
aaa authentication enable default group RADIUS-SERVERS enable
aaa authorization exec default group RADIUS-SERVERS if-authenticated
aaa authorization network default group RADIUS-SERVERS if-authenticated
aaa accounting send stop-record authentication failure
aaa session-id common
!
line con 0
logging synchronous
stopbits 1
line vty 0 4
logging synchronous
transport input ssh
```

Refer to the exhibit. A network administrator successfully logs in to a switch using SSH from a RADIUS server. When the network administrator uses a console port to access the switch, the RADIUS server returns shell:priv-lvl=15" and the switch asks to enter the enable command. When the command is entered, it gets rejected.
Which command set is used to troubleshoot and resolve this issue?

A. line con 0 aaa authorization console privl5 ! line vty 0 4 authorization exec

B. line con 0 aaa authorization console ! line vty 0 4 authorization exec

C. line con 0 aaa authorization console authorization priv15 ! line vty 0 4 transport input ssh

D. line con 0 aaa authorization console authorization exec ! line vty 0 4 transport input ssh

Correct Answer: *D*
Reference:
https://flylib.com/books/en/1.233.1.74/1/

Community vote distribution
D (75%) B (25%)

ZamanR 5 days, 7 hours ago
D is correct
upvoted 1 times

Pietjeplukgeluk 3 weeks, 3 days ago
Almost sure it is D, but the command is a bit broken, D: "line con 0 aaa authorization console authorization exec ! line vty 0 4 transport input ssh "
It needs to be:
1. globally enable authorization on console: aaa authorization console
2. move to line console 0: line con 0
3. Set the group to be used for authorization (note default is missing in the answer): authorization exec default
4. Go to line vty 0 4 (will set the same twice on next step): line vty 0 4
5. setting transport again to ssh: transport input ssh

Note that "console authentication == DISABLED by default": more info https://flylib.com/books/en/1.233.1.74/1/

So concluding, answer is bad quality, but D seems best of them.
upvoted 1 times

Ll123123 1 month, 3 weeks ago
Selected Answer: B

I actually prefer B.
SSH has no problem login, so the authorisation for vty must work. B has vty authorisation exec which is the default authorisation rule, and console authentication should work already, so just need to enable aaa authorisation console, and line console 0 thus can be empty configured

upvoted 1 times

⊟ **inteldarvid** 5 months, 2 weeks ago

Selected Answer: D

option D:
https://itexamanswers.net/question/refer-to-the-exhibit-a-network-administrator-successfully-logs-in-to-a-switch-using-ssh-from-a-radius-server-when-the-network-administrator-uses-a-console-port-to-access-the-switch-the-radius-server

upvoted 1 times

⊟ **HungarianDish** 7 months, 3 weeks ago

"aaa authorization console" is a global command, so we won't apply it under the line configuration.
"authorization exec" is only a partial command combiened with an authorization list (global).
D is closest.

upvoted 4 times

⊟ **Titini** 10 months ago

Selected Answer: D

We need to enable aaa auth console and auth exec for console and D has them. I do not understand why the vty conf is repeated in D but is the only answer that resolves the issue.

upvoted 2 times

⊟ **VergilP** 1 year, 1 month ago

can anyone explain this?

upvoted 1 times

⊟ **jarz** 1 year, 1 month ago

I think the ans is B

upvoted 1 times

⊟ **jarz** 1 year, 1 month ago

I had to Lab this to understand it.

Of the answers provided, none are correct!

aaa commands aren't supported directly on the lines and that for this scenario to work the Global Command aaa authorization console needed to be added to the configuration!

upvoted 4 times

⊟ **VergilP** 1 year, 1 month ago

300-410 ENARSI have many confuse question for me ....
oh my god

upvoted 5 times

Question #162 Topic 1

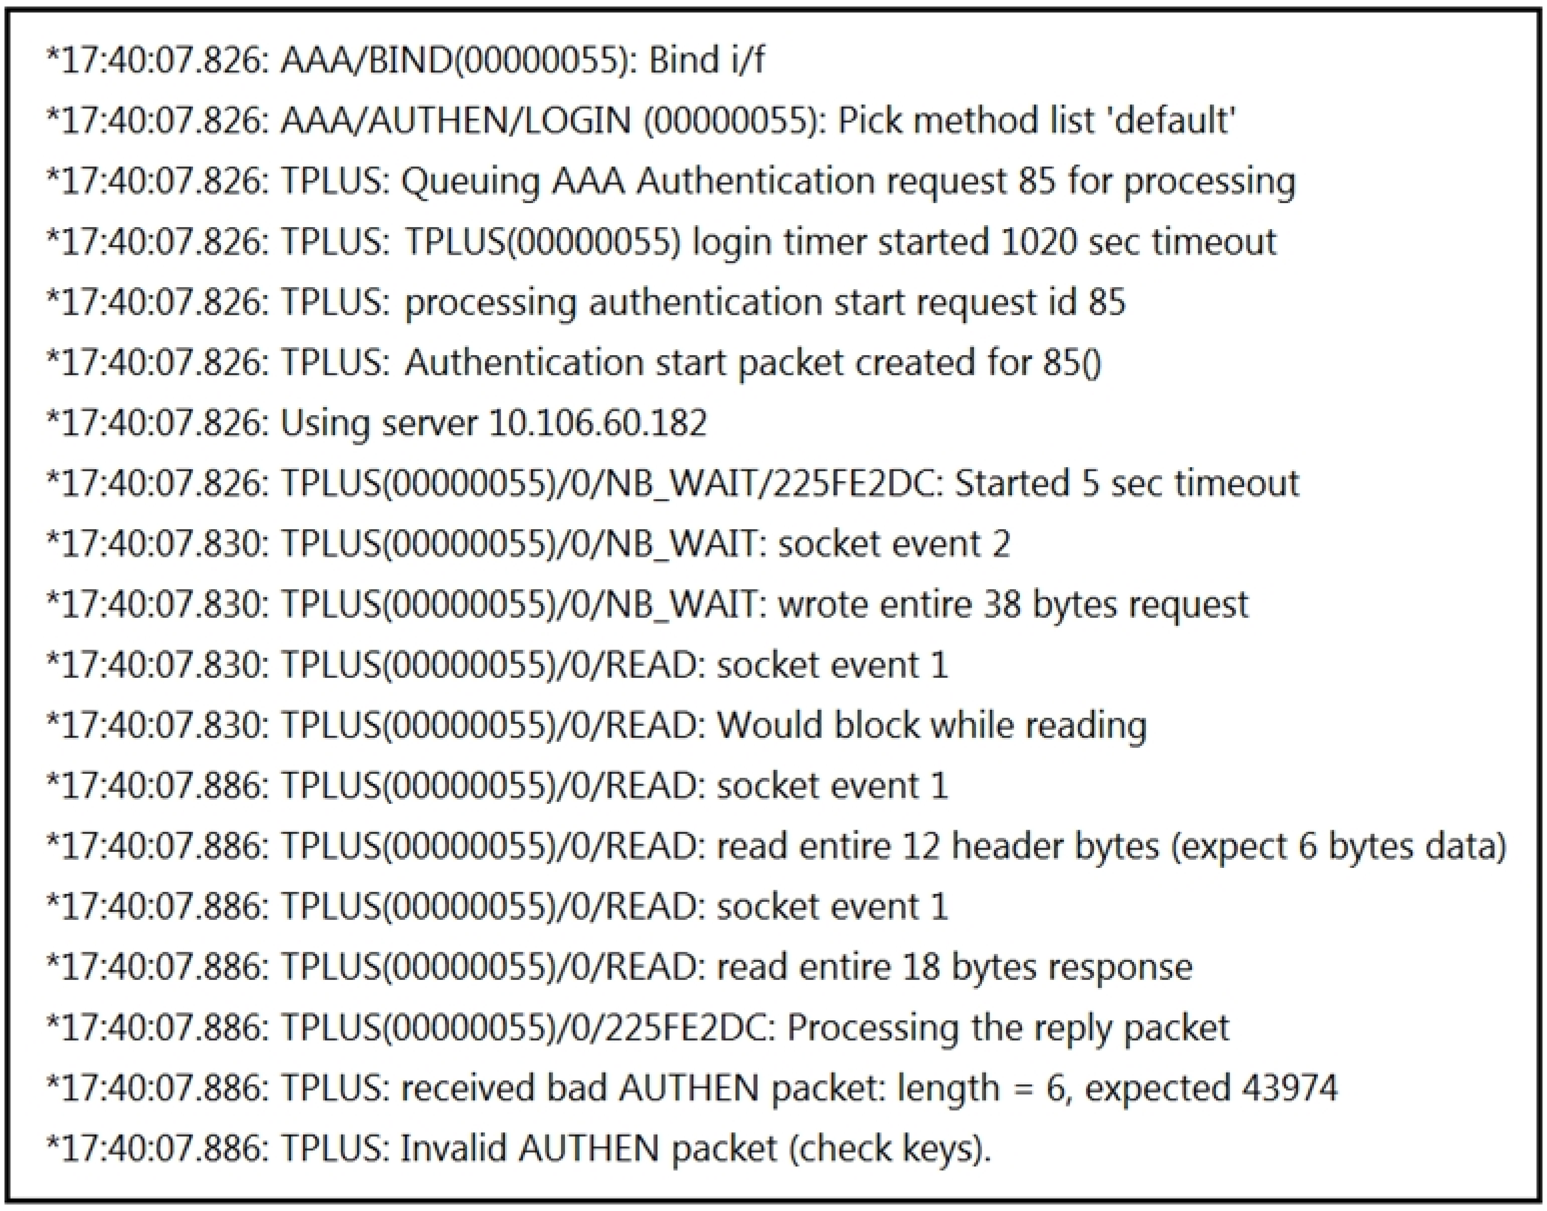

```
*17:40:07.826: AAA/BIND(00000055): Bind i/f
*17:40:07.826: AAA/AUTHEN/LOGIN (00000055): Pick method list 'default'
*17:40:07.826: TPLUS: Queuing AAA Authentication request 85 for processing
*17:40:07.826: TPLUS: TPLUS(00000055) login timer started 1020 sec timeout
*17:40:07.826: TPLUS: processing authentication start request id 85
*17:40:07.826: TPLUS: Authentication start packet created for 85()
*17:40:07.826: Using server 10.106.60.182
*17:40:07.826: TPLUS(00000055)/0/NB_WAIT/225FE2DC: Started 5 sec timeout
*17:40:07.830: TPLUS(00000055)/0/NB_WAIT: socket event 2
*17:40:07.830: TPLUS(00000055)/0/NB_WAIT: wrote entire 38 bytes request
*17:40:07.830: TPLUS(00000055)/0/READ: socket event 1
*17:40:07.830: TPLUS(00000055)/0/READ: Would block while reading
*17:40:07.886: TPLUS(00000055)/0/READ: socket event 1
*17:40:07.886: TPLUS(00000055)/0/READ: read entire 12 header bytes (expect 6 bytes data)
*17:40:07.886: TPLUS(00000055)/0/READ: socket event 1
*17:40:07.886: TPLUS(00000055)/0/READ: read entire 18 bytes response
*17:40:07.886: TPLUS(00000055)/0/225FE2DC: Processing the reply packet
*17:40:07.886: TPLUS: received bad AUTHEN packet: length = 6, expected 43974
*17:40:07.886: TPLUS: Invalid AUTHEN packet (check keys).
```

Refer to the exhibit. An engineer is troubleshooting a TACACS problem.
Which action resolves the issue?

A. Configure a matching TACACS server IP.

B. Configure a matching preshared key.

C. Generate authentication from a relative source interface.

D. Apply a configured AAA profile to the VTY.

Correct Answer: *B*
Reference:
https://community.cisco.com/t5/network-access-control/issues-with-tacacs-authentication/td-p/3412001

Community vote distribution

B (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

Look the keyword is "check key". Option is B

upvoted 3 times

HungarianDish 7 months, 3 weeks ago

Selected Answer: B

https://community.cisco.com/t5/network-access-control/bad-invalid-authentication-packet/td-p/824682

upvoted 3 times

Question #163 *Topic 1*

The network administrator configured CoPP so that all HTTP and HTTPS traffic from the administrator device located at 172.16 1.99 toward the router CPU is limited to 500 kbps. Any traffic that exceeds this limit must be dropped. access-list 100 permit ip host 172.16.1.99 any

```
!
class-map CM-ADMIN
match access-group 100
!
policy-map PM-COPP
class CM-ADMIN
police 500000 conform-action transmit
!
interface E0/0
service-policy input PM-COPP
```

CoPP failed to capture the desired traffic and the CPU load is getting higher.

Which two configurations resolve the issue? (Choose two.)

A. interface E0/0 no service-policy input PM-COPP ! control-plane service-policy input PM-COPP

B. policy-map PM-COPP class CM-ADMIN no police 500000 conform-action transmit police 500 conform-action transmit ! control-plane service-policy input PM-COPP

C. no access-list 100 access-list 100 permit tcp host 172.16.1.99 any eq 80

D. no access-list 100 access-list 100 permit tcp host 172.16.1.99 any eq 80 access-list 100 permit tcp host 172.16.1.99 any eq 443

E. policy-map PM-COPP class CM-ADMIN no police 500000 conform-action transmit police 500 conform-action transmit

Correct Answer: *A*

Community vote distribution

D (100%)

SAMAKEMM 2 months, 2 weeks ago

Correct answer: A & D

upvoted 1 times

[Removed] 5 months ago

Selected Answer: D

Choose two, A&D, configure the Control Plane to reference the Policy Map inbound. and the access list needs to reference por 80 and 443 for HTTP and HTTPS respectively.

upvoted 2 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

are two option: A and D

upvoted 1 times

Dacusai 7 months, 2 weeks ago

A&D, question says choose 2. But to restrict traffic to 500 kb we need to add the exceed-action drop command in order to do real control

upvoted 2 times

forccnp 9 months, 3 weeks ago

Selected Answer: D

A&D are correct answers

upvoted 1 times

Xerath 10 months ago

The answer is: A & D.

upvoted 1 times

rogabor81 11 months, 3 weeks ago

I would say A an D as well. but should not we add an exceed-action drop at the and as well? it says that any exceeding traffic should be dropped....

upvoted 4 times

HungarianDish 7 months, 3 weeks ago

That seems to be missing, too. A+D, plus exceed-action drop.
upvoted 1 times

Muste 4 months, 1 week ago
The default policing action if you only configured conformed-action is to drop the packets that exceed the configured rate limit.
upvoted 1 times

Noproblem22 1 year ago
A D are correct answer
upvoted 1 times

ChillingAgain 1 year, 1 month ago
A and D are correct.

Please correct the answers!
upvoted 1 times

xziomal9 1 year, 7 months ago
The correct answer is: A D
upvoted 1 times

Hack4 1 year, 7 months ago
A AND D
upvoted 1 times

piojo 1 year, 7 months ago
A and D (choose two)
upvoted 1 times

Question #164 Topic 1

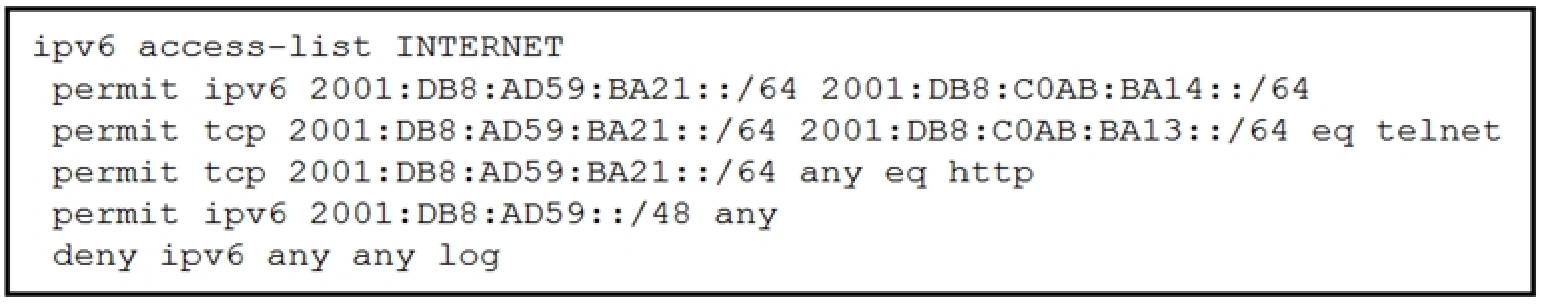

```
ipv6 access-list INTERNET
 permit ipv6 2001:DB8:AD59:BA21::/64 2001:DB8:C0AB:BA14::/64
 permit tcp 2001:DB8:AD59:BA21::/64 2001:DB8:C0AB:BA13::/64 eq telnet
 permit tcp 2001:DB8:AD59:BA21::/64 any eq http
 permit ipv6 2001:DB8:AD59::/48 any
 deny ipv6 any any log
```

Refer to the exhibit. While monitoring VTY access to a router, an engineer notices that the router does not have any filter and anyone can access the router with username and password even though an ACL is configured.
Which command resolves this issue?

A. access-class INTERNET in

B. ip access-group INTERNET in

C. ipv6 traffic-filter INTERNET in

D. ipv6 access-class INTERNET in

Correct Answer: *D*

Community vote distribution

D (85%) C (15%)

⊟ **TECH3K3** Highly Voted 1 year, 5 months ago

Selected Answer: D

Answer is D:

IPv6 access-class vs IPv6 traffic-filter
The difference depends on whether you want to filter IPv6 traffic sent *to* the router or *through* the router.

The 'ipv6 traffic-filter' command is used to filter IPv6 traffic flowing through an interface:
Command reference (with example):
http://www.cisco.com/en/US/docs/ios/ipv6/command/reference/ipv6_09.html#wp2297000

The 'ipv6 access-class' command is used to filter IPv6 traffic destined to the router (i.e. management traffic).
Command reference (with example):
http://www.cisco.com/en/US/docs/ios/ipv6/command/reference/ipv6_05.html#wp2274594

upvoted 12 times

⊟ **asans** Most Recent 3 days, 17 hours ago

Selected Answer: C

Both C and D works to filter telnet access but in this case the acl, INTERNET, is not only dealing with telnet traffic but http and hosts as well and so it has to be applied at the interface using ipv6 traffic-filter in. C is the correct answer

upvoted 1 times

⊟ **asans** 3 days, 17 hours ago

Both C and D works to filter telnet access but in this case the acl, INTERNET, is not only dealing with telnet traffic but http and hosts as well and so it has to be applied at the interface using ipv6 traffic-filter in. C is the correct answer

upvoted 1 times

⊟ **Wh00py** 4 months ago

Answer is D:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_acl/configuration/xe-16/sec-data-acl-xe-16-book/ip6-acls-xe.html

upvoted 1 times

⊟ **Cyril_the_Squirl** 4 months, 1 week ago

How can so many people get it wrong?
traffic-filter command is the ipv6 equivalent for ip access-group for applying access-list to an interface

upvoted 1 times

⊟ **Slinky** 8 months, 2 weeks ago

This is being applied to the vty lines, so the answer is D

upvoted 1 times

⊟ **chikuwan** 1 year, 4 months ago

Selected Answer: D

first, you should define ipv6 access-list in grobal configuration mode,and ipv6 traffic-filter is when you want to apply it in a interface, and when in conditio of a vty ,the command wull be access-list, the answer is D,given answer is correct

upvoted 2 times

⊟ **Nhan** 1 year, 6 months ago

C is correct answer, the ipv6 access-list need to be applied on an interface using ip filter command

upvoted 1 times

⊟ **piojo** 1 year, 6 months ago

Selected Answer: D

Simulated in a lab.

It also can be applied to the vty with ipv6 access-class command.

So, examine if the access-list applied via ipv6 access-class permit tcp traffic to port 23 (or 22 when ssh) from / to the desired IPs.

upvoted 3 times

⊟ **timtgh** 1 year, 6 months ago

C is right,

upvoted 1 times

⊟ **Kimaf** 1 year, 8 months ago

Selected Answer: C

This is the right command to apply to the interface.

upvoted 2 times

Question #165 | *Topic 1*

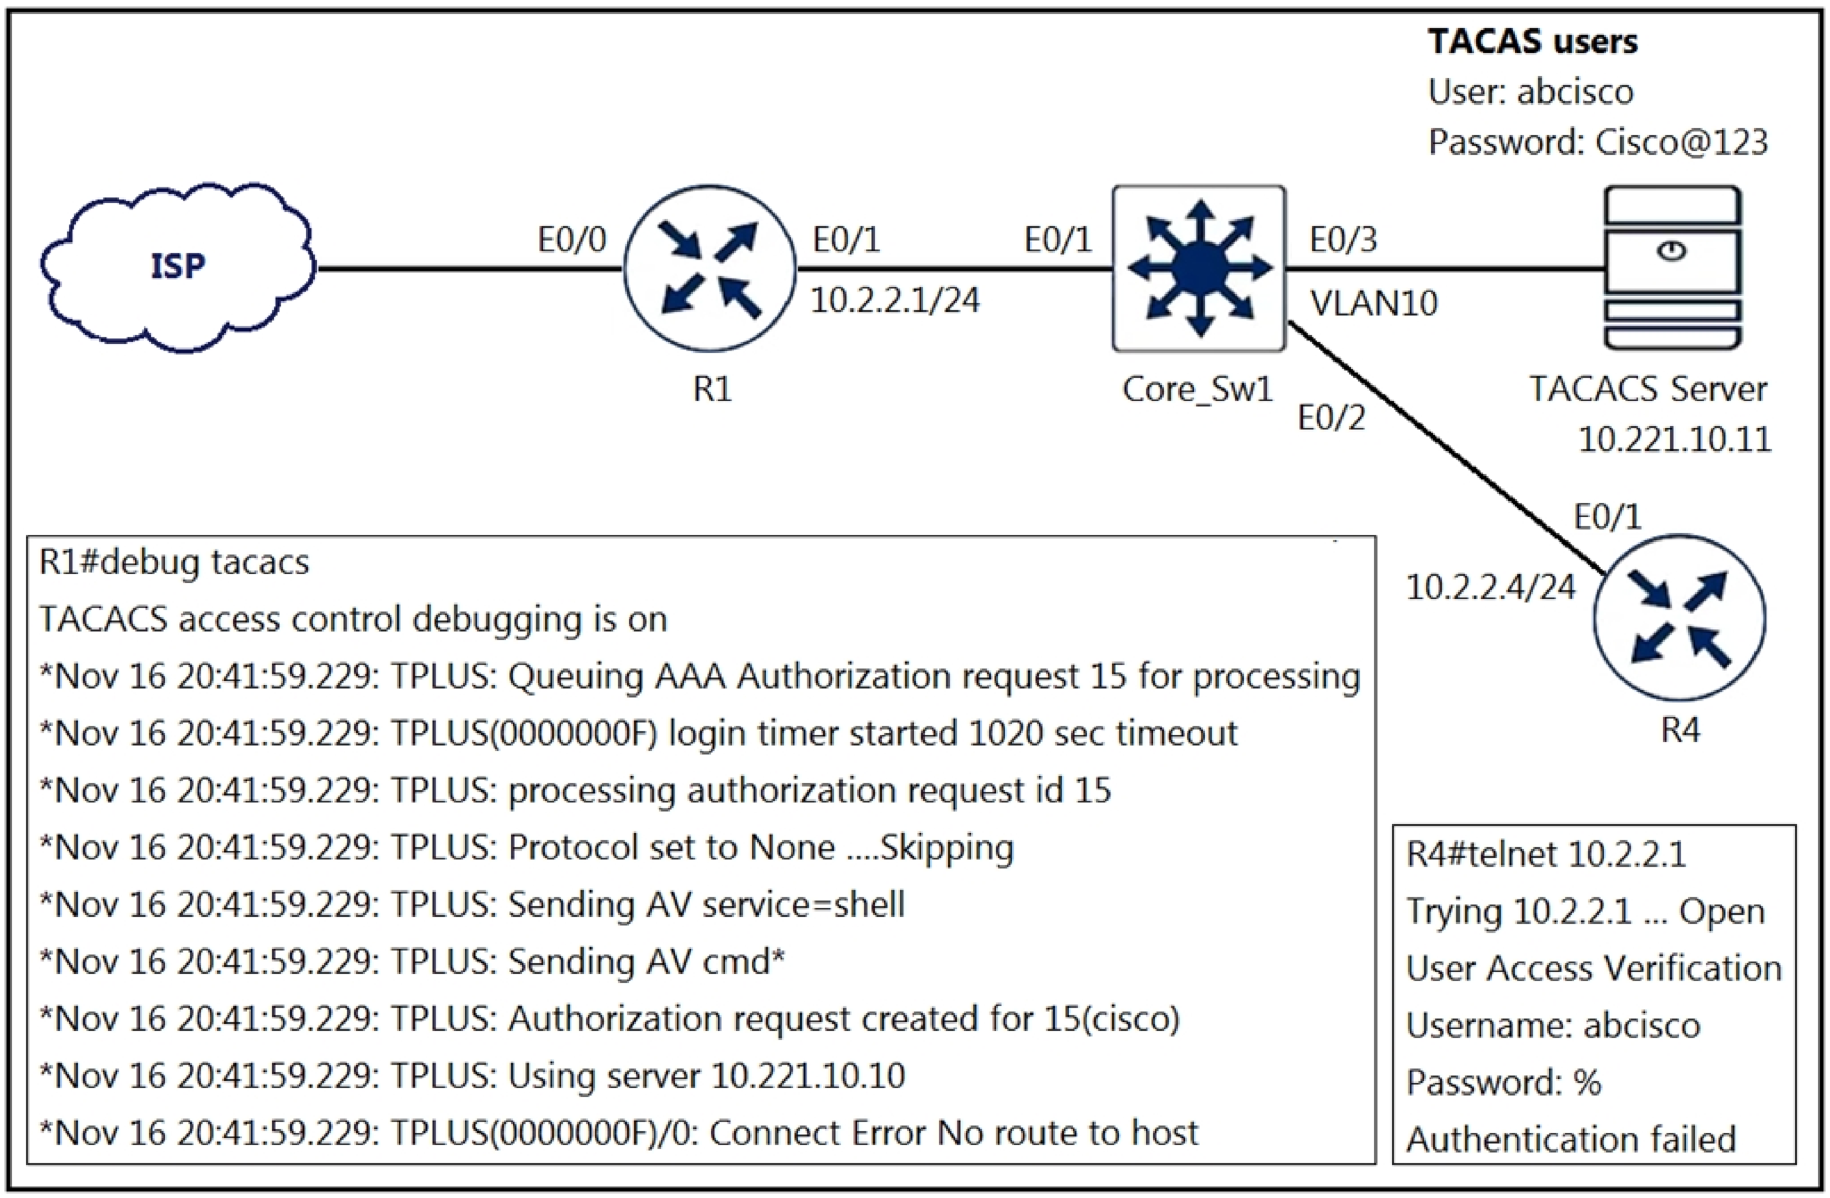

```
R1#debug tacacs
TACACS access control debugging is on
*Nov 16 20:41:59.229: TPLUS: Queuing AAA Authorization request 15 for processing
*Nov 16 20:41:59.229: TPLUS(0000000F) login timer started 1020 sec timeout
*Nov 16 20:41:59.229: TPLUS: processing authorization request id 15
*Nov 16 20:41:59.229: TPLUS: Protocol set to None ....Skipping
*Nov 16 20:41:59.229: TPLUS: Sending AV service=shell
*Nov 16 20:41:59.229: TPLUS: Sending AV cmd*
*Nov 16 20:41:59.229: TPLUS: Authorization request created for 15(cisco)
*Nov 16 20:41:59.229: TPLUS: Using server 10.221.10.10
*Nov 16 20:41:59.229: TPLUS(0000000F)/0: Connect Error No route to host
```

```
R4#telnet 10.2.2.1
Trying 10.2.2.1 ... Open
User Access Verification
Username: abcisco
Password: %
Authentication failed
```

Refer to the exhibit. An engineer is trying to connect to R1 via Telnet with no success.
Which configuration resolves the issue?

A. tacacs server prod address ipv4 10.221.10.10 exit

B. ip route 10.221.10.10 255.255.255.255 ethernet 0/1

C. ip route 10.221.0.11 255.255.255.255 ethernet 0/1

D. tacacs server prod address ipv4 10.221.10.11 exit

Correct Answer: *C*

Community vote distribution

D (92%) 4%

JingleJangus Highly Voted 1 year, 7 months ago

Selected Answer: D

No one is going to say anything about this one?
From what I can tell, C really isn't the BEST answer for this question.
In my opinion, D is a much better answer.

Reason:
Looking at the debug output, second to the last line, the log suggests that it is attempting to use server x.x.10.10 when the diagram specifies that the server is actually x.x.10.11.
This would require the tacacs group to be modified to use the correct server:
`tacacs server prod`
`address ipv4 x.x.10.11`
`exit`

I suppose C is a good backup answer, but D ensures that we are pointing to the correct tacacs server to begin with.

upvoted 14 times

inteldarvid Most Recent 5 months, 2 weeks ago

Selected Answer: D

D correct

upvoted 1 times

Dacusai 7 months, 3 weeks ago

I just hope that the exam has the correct questions and answers

upvoted 4 times

Calyfas 9 months, 3 weeks ago

Selected Answer: D

I believe that is option D, in the diagram the tacacs server has ip address 10.221.10.11. So you fix the ip address in router config. There is no other TACACS server in place. So, option B is wrong.

upvoted 2 times

tseen 10 months, 2 weeks ago

Selected Answer: D

Second to last line of the debug shows that the wrong TACACS server IP address of 10.221.10.10 was configured instead of the correct TACACS server IP address 10.221.10.11. Hence configuring the correct TACACS server IP address(10.221.10.11) will solve the problem

upvoted 1 times

NoUserName1234 1 year ago

Selected Answer: B

For the question it needs to be Answer B too fix the complete issue it would be B&D

upvoted 1 times

NoUserName1234 1 year ago

Selected Answer: C

The output suggest that there is no route too the tacacs server. Soo the correct answer would be C because you need to set a route. That the wrong Tacacs server is used is another issue

upvoted 1 times

TECH3K3 1 year, 5 months ago

Selected Answer: D

The ip address of the server being used is wrong

upvoted 2 times

Sunsammie 1 year, 5 months ago

C is the answer as there is no route to the tacacs server. How will the debug know of the server address if it had not been configured.

upvoted 1 times

Nhan 1 year, 5 months ago

After careful ready the question I would like to withdraw my previous statement, the given answer is correct, there is no route to the host that causing the authentication fail, so ac is the right answer

upvoted 1 times

phryde 1 year, 6 months ago

Selected Answer: D

D is the only one correcting the IP of the TACACS server

upvoted 2 times

Nhan 1 year, 6 months ago

this is authentication issue, setting a static route having nothing to do with authentication, therefore C is not a good answer even its help the establish the route to the device directly,
again, since this is authentication issue, pointing to the correct AAA server is more important so D would be correct answer

upvoted 1 times

xziomal9 1 year, 7 months ago

Selected Answer: D

The correct answer is: D

upvoted 1 times

len184 1 year, 7 months ago

I select D because host is not using the server as specified in the diagram.

upvoted 2 times

Question #166 *Topic 1*

An engineer is trying to copy an IOS file from one router to another router by using TFTP.
Which two actions are needed to allow the file to copy? (Choose two.)

A. Copy the file to the destination router with the copy tftp: flash: command

B. Enable the TFTP server on the source router with the tftp-server flash: <filename> command

C. TFTP is not supported in recent IOS versions, so an alternative method must be used

D. Configure a user on the source router with the username tftp password tftp command

E. Configure the TFTP authentication on the source router with the tftp-server authentication local command

Correct Answer: *AB*

Community vote distribution

AB (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: AB

A and B correct

upvoted 1 times

Noproblem22 1 year ago

AB are correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 3 times

Question #167 *Topic 1*

Refer to the exhibit. Users report that IP addresses cannot be acquired from the DHCP server. The DHCP server is configured as shown. About 300 total nonconcurrent users are using this DHCP server, but none of them are active for more than two hours per day.
Which action fixes the issue within the current resources?

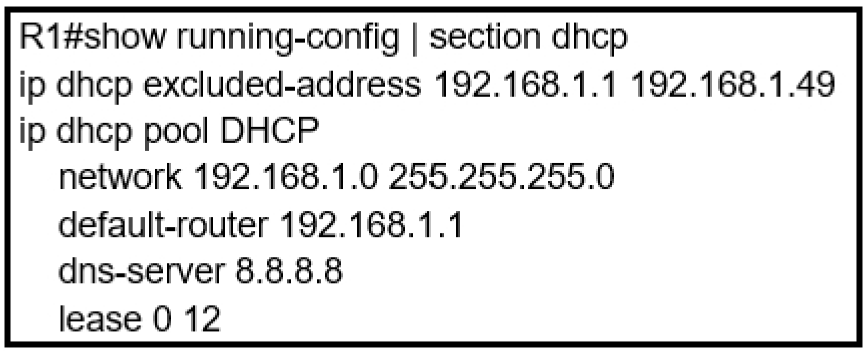

```
R1#show running-config | section dhcp
ip dhcp excluded-address 192.168.1.1 192.168.1.49
ip dhcp pool DHCP
   network 192.168.1.0 255.255.255.0
   default-router 192.168.1.1
   dns-server 8.8.8.8
   lease 0 12
```

A. Modify the subnet mask to the network 192.168.1.0 255.255.254.0 command in the DHCP pool

B. Configure the DHCP lease time to a smaller value

C. Configure the DHCP lease time to a bigger value

D. Add the network 192.168.2.0 255.255.255.0 command to the DHCP pool

Correct Answer: *B*

Community vote distribution

B (100%)

Mohammad963 4 months, 1 week ago

B is Correct, as the users, nonconcurrent

upvoted 1 times

Chiaretta 5 months, 1 week ago

Selected Answer: B

B is correct because it says that ip adresses are used for two hour per day

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

option B is corerct, because its necesary for all user in 2 hours p-day

upvoted 1 times

JoeyT 7 months, 2 weeks ago

why A wrong? For B, if we have 290 concurrent users which satisfy the question but still won't work

upvoted 1 times

[Removed] 4 months ago

I thought about this, but the last part of the question: "which action fixes the issue WITHIN THE CURRENT RESOURCES" gave the hint that we are not allowed to increase the address space.
B is the best answer.

upvoted 1 times

Calyfas 9 months, 3 weeks ago

Selected Answer: B

Option B is the only that makes sense to me.

upvoted 1 times

Ckl22 1 year ago

Selected Answer: B

Which action fixes the issue within the current resources?

By changing the lease time, it doesn't require an increase in resources

upvoted 1 times

Noproblem22 1 year ago

B makes sense

upvoted 1 times

TECH3K3 1 year, 4 months ago

Selected Answer: B

This isn't rocket science.
Each user only needs a DHCP IP for no longer than 2 hours.
Currently, the lease time is 12 hours, so for an average of 10 hours IP addresses are tied up and not release back into the DHCP pool
I would configure a lease time of 1-1.5 hour. If the client still needs an IP address, then they will be issued back the same IP it was originally using.

upvoted 2 times

ellen_AA 11 months, 3 weeks ago

It always depends. Sometimes, like in a home network. A lease of 2 days is alright.

upvoted 1 times

Jacklee2022 1 year, 10 months ago

If have got 204-299 users are concurrent active, so What is happening? C or D is correct in this case?

upvoted 1 times

Jacklee2022 1 year, 10 months ago

In here answers are change sorting, my mind is A or B in this question

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

TECH3K3 1 year, 5 months ago

ExamShark, this is all you say in very comment and NEVER once explain why you agree or add any value.

upvoted 4 times

uglyprawn 2 years, 9 months ago

why not incresing the host portion? it will allow 510 host?

upvoted 4 times

jjj554 2 years, 9 months ago

that would solve the problem only if we also alter the interface hosting the network, but theres no mention of altering that so I assume we go with the one that will still solve the issue without additional configurations.

upvoted 2 times

Macferson 1 year, 10 months ago

the trick here is the number of 300 users, but it is not the real problem since they are not concurrent this means that they are not disputing leased IPs. The issue here is the lease time so it needs to be changed, that is why the answer is B.

upvoted 1 times

akbntc 3 years ago

The keywords in the question are:
1. 300 non-concurrent users
2. They connect only 2 hours per day
So, decreasing the lease time definitely solves the problem. B is correct.

upvoted 3 times

anonymous1966 3 years, 5 months ago

B is correct. DHCP server should release the addresses faster for the other clients.

upvoted 2 times

heamgu 3 years, 5 months ago

Correct answer is C.

lease [Days][Hours][Minutes]

upvoted 4 times

Earl03 3 years, 5 months ago

That is wrong.
You have a total of 203 IP addresses, and 300 clients. At the current configuration every Client up to the 203rd gets an IP, and occupies it for 12 hours (0 days 12 hours 0 minutes).
So the 204th client per day fails to get an IP address.
To fix this we need to lower the IP lease time from 12 hours down to somewhere around 2 hours as suggested in the question.

upvoted 7 times

Refer to the exhibit. ISP 1 and ISP 2 directly connect to the Internet. A customer is tracking both ISP links to achieve redundancy and cannot see the Cisco IOS IP

SLA tracking output on the router console.

Which command is missing from the IP SLA configuration?

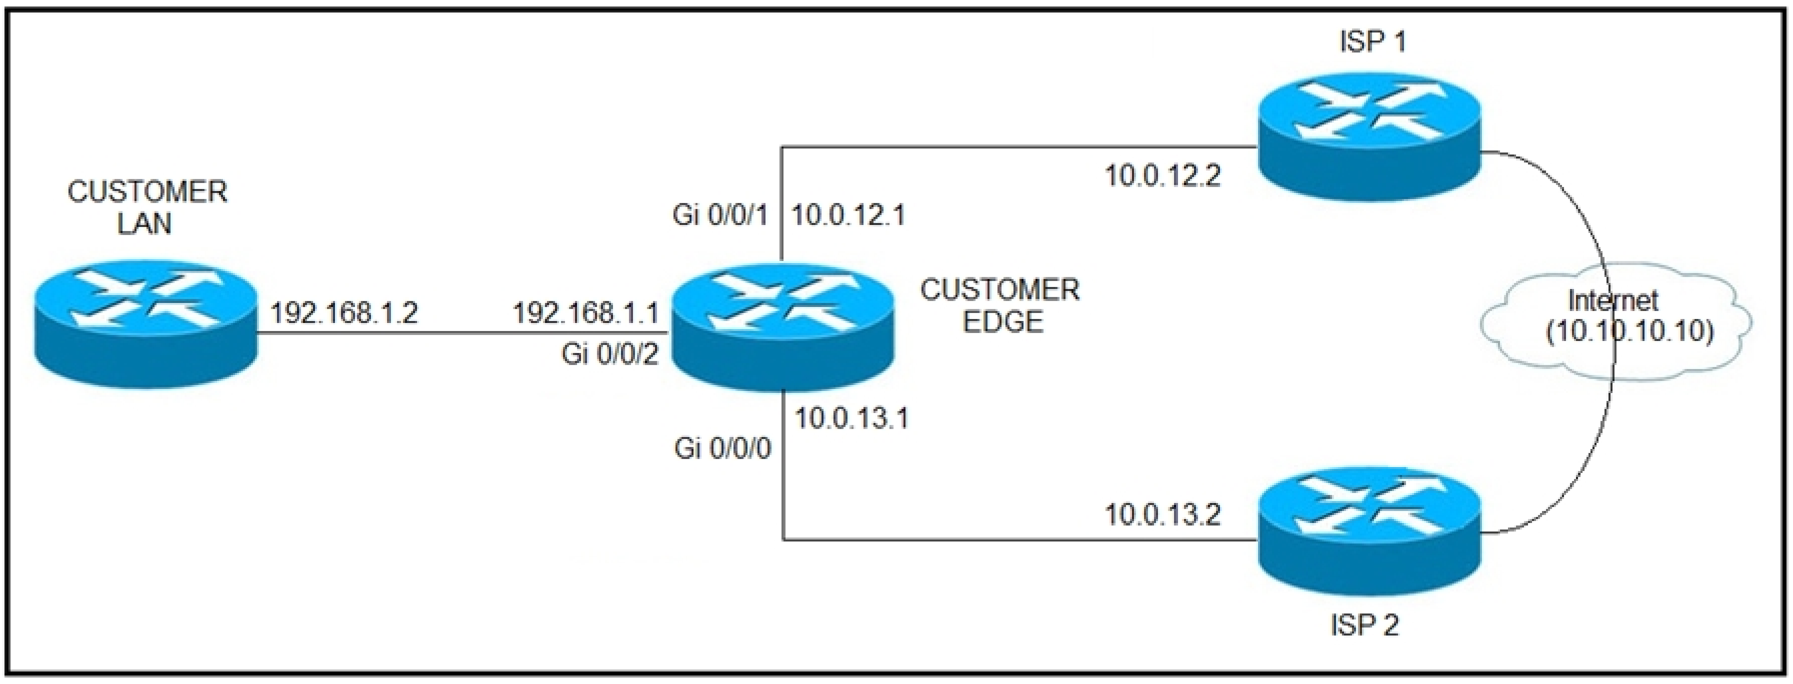

A. Start-time 00:00

B. Start-time 0

C. Start-time immediately

D. Start-time now

Correct Answer: *D*

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/configuration/15-mt/sla-15-mt-book/sla_icmp_echo.html

Community vote distribution

D (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

Option correct is D:
Customer needs to run it as soon as possible
ip sla schedule 1 life forever start-time now

upvoted 1 times

HungarianDish 7 months, 3 weeks ago

track 1 ip sla 1 reachability
ip sla 1
icmp-echo <target IP>
ip sla schedule 1 life forever start-time now
-and if designed like that, then they might add the track statement to the static default route pointing to one ISP, and make the static default route to the other ISP a floating static route
https://www.firewall.cx/cisco-technical-knowledgebase/cisco-routers/813-cisco-router-ipsla-basic.html

upvoted 1 times

Kimaf 1 year, 8 months ago

This question is missing configuration

upvoted 3 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

Question #169 Topic 1

Refer to the exhibit. An administrator noticed that after a change was made on R1, the timestamps on the system logs did not match the clock.

What is the reason for this error?

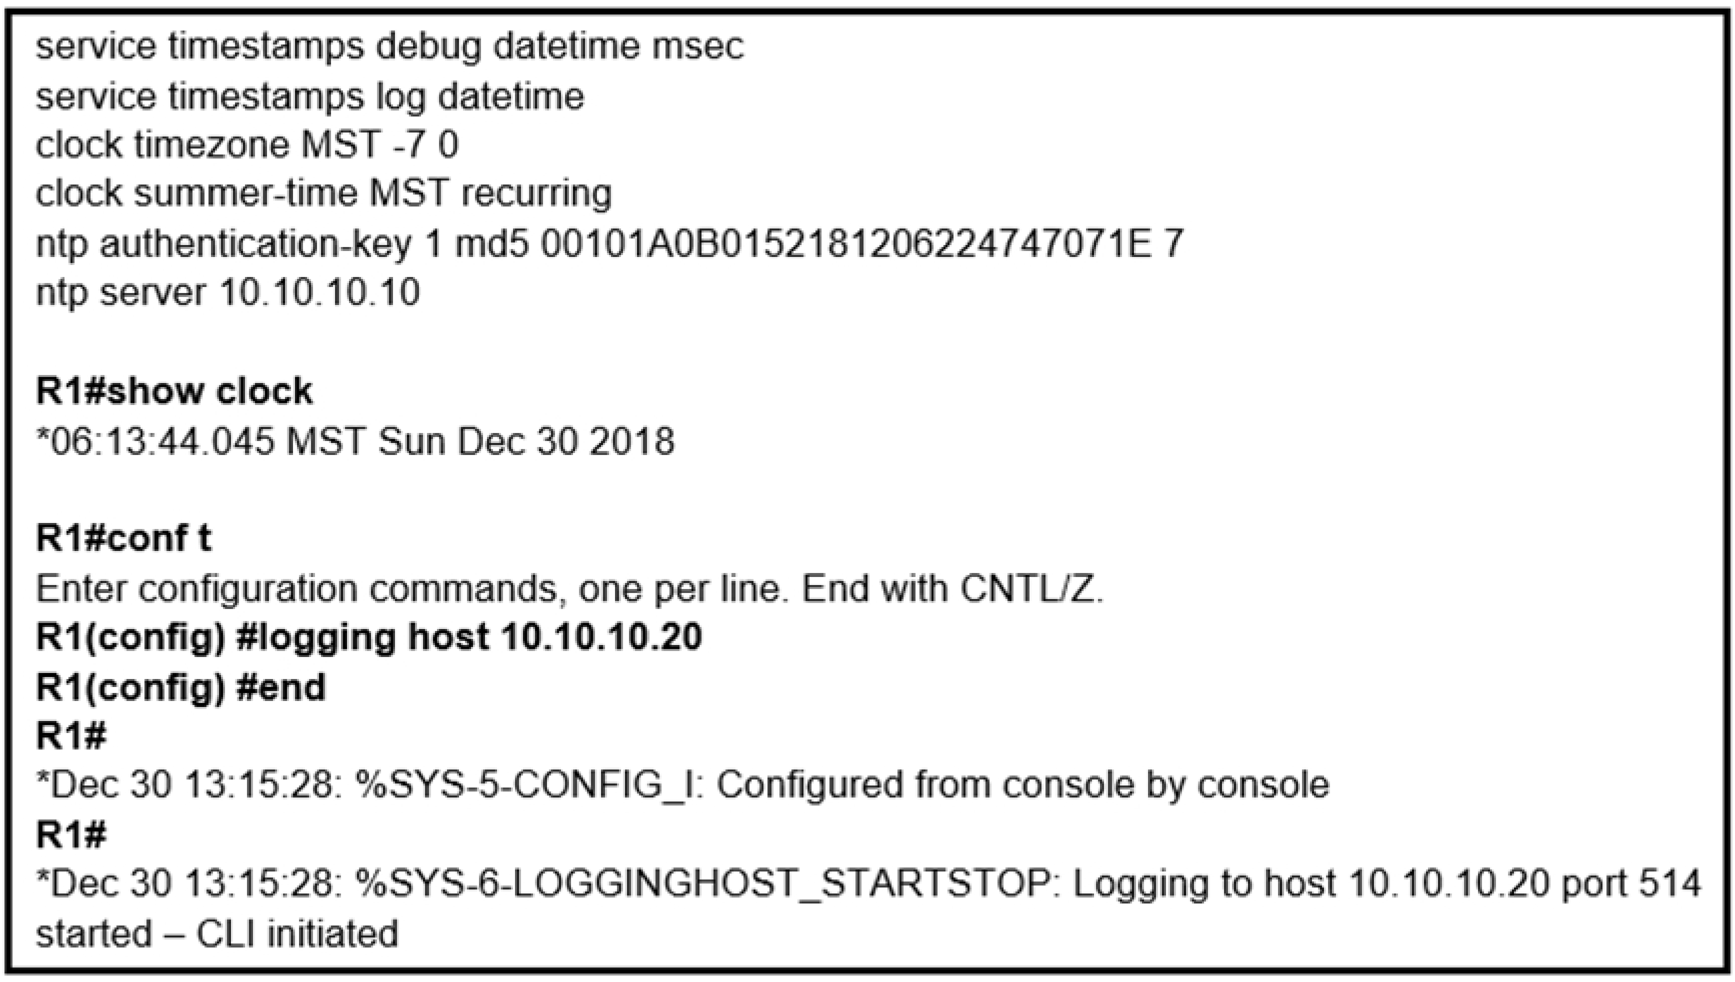

```
service timestamps debug datetime msec
service timestamps log datetime
clock timezone MST -7 0
clock summer-time MST recurring
ntp authentication-key 1 md5 00101A0B0152181206224747071E 7
ntp server 10.10.10.10

R1#show clock
*06:13:44.045 MST Sun Dec 30 2018

R1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config) #logging host 10.10.10.20
R1(config) #end
R1#
*Dec 30 13:15:28: %SYS-5-CONFIG_I: Configured from console by console
R1#
*Dec 30 13:15:28: %SYS-6-LOGGINGHOST_STARTSTOP: Logging to host 10.10.10.20 port 514
started – CLI initiated
```

A. An authentication error with the NTP server results in an incorrect timestamp.

B. The keyword localtime is not defined on the timestamp service command.

C. The NTP server is in a different time zone.

D. The system clock is set incorrectly to summer-time hours.

Correct Answer: *A*

Community vote distribution

B (64%) A (36%)

LuigiG Highly Voted 3 years, 6 months ago

I think B is the correct
https://community.cisco.com/t5/networking-documents/router-log-timestamp-entries-are-different-from-the-system-clock/ta-p/3132258

upvoted 20 times

Colmenarez Most Recent 3 months, 3 weeks ago

Selected Answer: A

It's matching that they are showing you the #show clock command. But the NTP is not working properly due "recent changes" you can notice that with the "*" symbol.

upvoted 1 times

Pietjeplukgeluk 3 weeks, 3 days ago

The question is why the clock is having a different time than the actual logs generated. The question is not that the clock is having the incorrect time. Concluding, the answer to the question appears to be B

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

yes option B
https://community.cisco.com/t5/networking-documents/router-log-timestamp-entries-are-different-from-the-system-clock/ta-p/3132258

upvoted 1 times

Rob_CCNP000 6 months ago

Selected Answer: A

"*" symbol means time is not authoritative: the software clock is not in sync or has never been set.

upvoted 1 times

Edwinmolinab 1 year, 1 month ago

Selected Answer: A

* at start shows a clock error against ntp server

upvoted 2 times

Hack4 1 year, 10 months ago

B is correct

upvoted 1 times

weltongama 1 year, 10 months ago

Selected Answer: B

B is the correct answer!

upvoted 2 times

tyh391 1 year, 10 months ago

Selected Answer: B

As in discussion

upvoted 2 times

myrmike 1 year, 11 months ago

The answer is B. Even if the the time was synched with NTP the log and the clock would still show a difference because localtime was not included in the service timestamp command

upvoted 3 times

spamke 2 years ago

Selected Answer: B

the questionnis why logs and clock do not match and not if ntp is syncd or not

upvoted 2 times

spamke 2 years ago

so it's B

upvoted 1 times

myrmike 2 years ago

A is correct. An '*' preceding the datetime indicates that the local clock is not synced with a time source

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is B

upvoted 2 times

Wesgo 2 years, 8 months ago

It's B hands down

upvoted 1 times

Malataw 3 years ago

By default, syslog and debug messages are stamped by UTC, regardless of the time zone that device configured. You should append localtime key word to "service timestamp {log | debug} datetime msec" global command to change that behavior.

upvoted 4 times

anonymous1966 3 years, 5 months ago

B is correct.

upvoted 3 times

heamgu 3 years, 5 months ago

Agree on B is the correct.

upvoted 3 times

Question #170 *Topic 1*

DRAG DROP -

Drag and drop the DHCP messages from the left onto the correct uses on the right.

Select and Place:

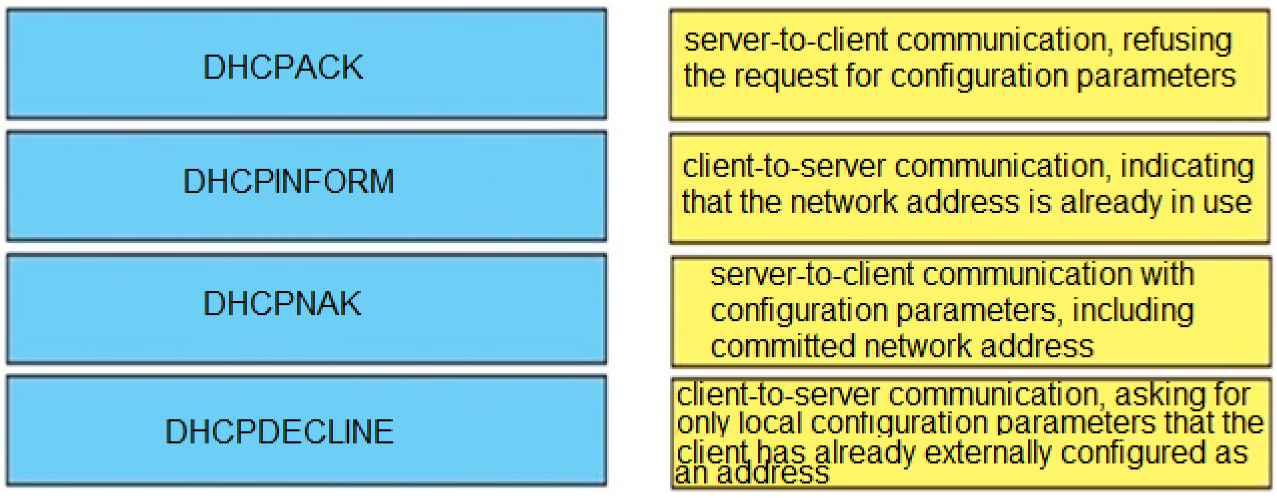

DHCPACK	server-to-client communication, refusing the request for configuration parameters
DHCPINFORM	client-to-server communication, indicating that the network address is already in use
DHCPNAK	server-to-client communication with configuration parameters, including committed network address
DHCPDECLINE	client-to-server communication, asking for only local configuration parameters that the client has already externally configured as an address

Correct Answer:

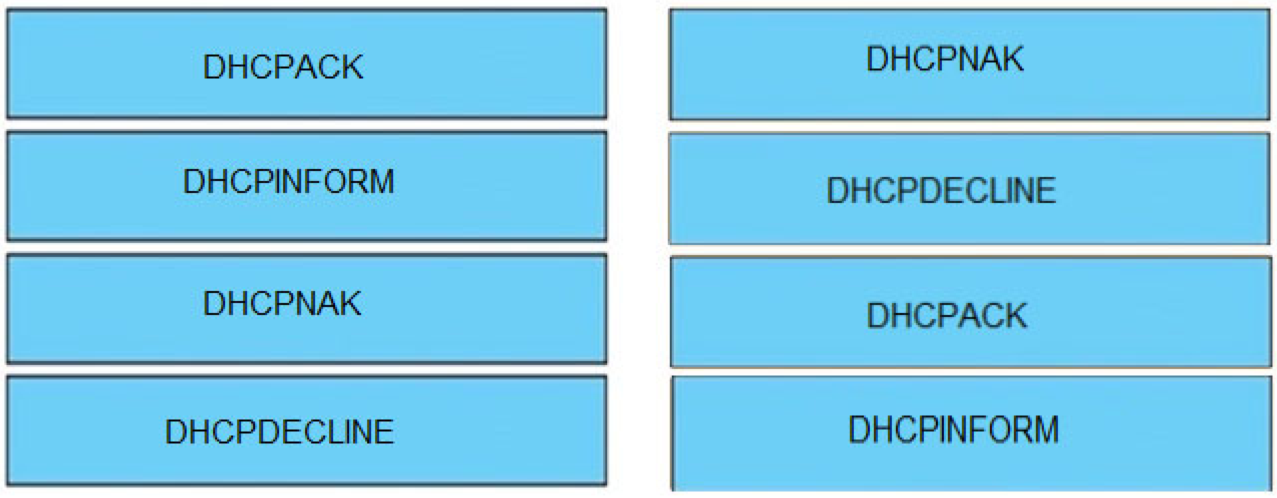

DHCPACK	DHCPNAK
DHCPINFORM	DHCPDECLINE
DHCPNAK	DHCPACK
DHCPDECLINE	DHCPINFORM

Reference:

https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html

Eric0_0 1 year, 9 months ago

https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html

upvoted 2 times

studybuddy10 2 years, 1 month ago

Agree too.

upvoted 1 times

Nik113 2 years, 1 month ago

totally agree @bjromeo

upvoted 1 times

bjromero28 2 years, 1 month ago

*DHCPNAK - Server to client negative acknowledgment indicating the client's understanding of the network address is incorrect (for example, if the client has moved to a new subnet), or a client's lease has expired.

*DHCPDECLINE - Client to server message indicating the network address is already being used.

*DHCPACK - Server to client acknowledgment message containing configuration parameters, including a confirmed network address.

*DHCPINFORM - Client to server message requesting only local configuration parameters; client has an externally configured network address.

Given Answer is correct

upvoted 3 times

Question #171 *Topic 1*

A network engineer is investigating a flapping (up/down) interface issue on a core switch that is synchronized to an NTP server. Log output currently does not show the time of the flap.

Which command allows the logging on the switch to show the time of the flap according to the clock on the device?

A. service timestamps log uptime

B. clock summer-time mst recurring 2 Sunday mar 2:00 1 Sunday nov 2:00

C. service timestamps log datetime localtime show-timezone

D. clock calendar-valid

Correct Answer: *C*

Community vote distribution

C (100%)

HungarianDish 7 months, 3 weeks ago

Selected Answer: C

We certainly need "service timestamps log ". The uptime of the switch is not relevant, so we do not need solution A). However, localtime is useful for troubleshooting.
https://www.cisco.com/E-Learning/bulk/public/tac/cim/cib/using_cisco_ios_software/cmdrefs/service_timestamps.htm
-enables time stamps on logging messages, showing the current time and date relative to the local time zone, with the time zone name included:
#service timestamps log datetime localtime show-timezone
log: Applies timestamps to logging messages.
localtime: Use local time zone for timestamps
show-timezone: Add time zone information to timestamp

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

Question #172 Topic 1

When provisioning a device in Cisco DNA Center, the engineer sees the error message `Cannot select the device. Not compatible with template`.
What is the reason for the error?

A. The template has an incorrect configuration.

B. The software version of the template is different from the software version of the device.

C. The changes to the template were not committed.

D. The tag that was used to filter the templates does not match the device tag.

Correct Answer: *D*
Reference:
https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/1-2-10/user_guide/b_cisco_dna_center_ug_1_2_10/b_dnac_ug_1_2_10_chapter_0111.html

Community vote distribution

D (100%)

Pizzadoos Highly Voted 2 years, 9 months ago
Just wanted to say that a question like this should not be part of the exam as this is what the exam topics list mentions for Cisco DNA center:
4.7 Troubleshoot network problems using Cisco DNA Center assurance (connectivity, monitoring, device health, network health)
upvoted 16 times

Pietjeplukgeluk 3 weeks, 1 day ago
Somehow Cisco thinks we need to learn Cisco errors and workflows in great detail. I cannot understand how this would lead to better engineers.
upvoted 1 times

Colmenarez Most Recent 3 months, 1 week ago
so, layer 8 issue.
upvoted 1 times

inteldarvid 5 months, 2 weeks ago
Selected Answer: D
D is correct:

https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/2-1-2/user_guide/b_cisco_dna_center_ug_2_1_2/b_cisco_dna_center_ug_2_1_1_chapter_01000.html

If you use tags to filter the templates, you must apply the same tags to the device to which you want to apply the templates. Otherwise, you get the following error during provisioning: "Cannot select the device. Not compatible with template."
upvoted 1 times

Kayyye 2 years ago
The given answer is correct
upvoted 1 times

examShark 2 years, 4 months ago
The given answer is correct
upvoted 3 times

thissiteisgreat 2 years, 11 months ago
Update the link:
https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/1-3-3-0/user_guide/b_cisco_dna_center_ug_1_3_3_0/b_cisco_dna_center_ug_1_3_2_0_chapter_01000.html?bookSearch=true

'If you use tags to filter the templates, you must apply the same tags to the device to which you want to apply the templates. Otherwise, the following error occurs during provisioning: "Cannot select the device. Not compatible with template."'

So, the answer is correct.
upvoted 3 times

Question #173 Topic 1

While working with software images, an engineer observes that Cisco DNA Center cannot upload its software image directly from the device. Why is the image not uploading?

A. The device must be resynced to Cisco DNA Center.

B. The software image for the device is in install mode.

C. The device has lost connectivity to Cisco DNA Center.

D. The software image for the device is in bundle mode

Correct Answer: *B*
Reference:
https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/1-2-10/user_guide/b_cisco_dna_center_ug_1_2_10/b_dnac_ug_1_2_10_chapter_0100.html

Kayyye 2 years ago

The given answer is correct

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

thissiteisgreat 2 years, 11 months ago

Updated the link:
https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/1-3-3-0/user_guide/b_cisco_dna_center_ug_1_3_3_0/b_cisco_dna_center_ug_1_3_2_0_chapter_0100.html?bookSearch=true#id_77074

"When a device is in Install Mode, Cisco DNA Center is unable to upload its software image directly from the device. When a device is in install mode"

So, the answer is correct

upvoted 4 times

Question #174 *Topic 1*

An engineer configured the wrong default gateway for the Cisco DNA Center enterprise interface during the install.
Which command must the engineer run to correct the configuration?

A. sudo maglev-config update

B. sudo maglev install config update

C. sudo maglev reinstall

D. sudo update config install

Correct Answer: *A*

Community vote distribution

A (100%)

examShark Highly Voted 2 years, 4 months ago

The given answer is correct

https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/2-1-2/install_guide/2ndGen/b_cisco_dna_center_install_guide_2_1_2_2ndGen/m_troubleshoot_deployment_2_1_2_2ndgen.html

upvoted 7 times

KaFi_PaOr Most Recent 5 months, 2 weeks ago

OMG what the question! On Guide press book nothing about config in DNA Center

upvoted 2 times

[Removed] 4 months ago

Literally the book even says so:
"Cisco DNA Center is a massive topic that is beyond the scope of the ENARSI exam. The official exam objectives for ENARSI state that you should be able to “troubleshoot network problems using Cisco DNA Center Assurance (connectivity, monitoring, device health, net□work health).” Therefore, this section remains focused on this objective."

Cisco has some horrible team making these exams.

upvoted 2 times

HungarianDish 7 months, 3 weeks ago

Selected Answer: A

https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/BRKSDN-1029.pdf
https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2020/pdf/BRKNMS-2426.pdf

upvoted 1 times

Kayyye 2 years ago

The given answer is correct

upvoted 1 times

Question #175 *Topic 1*

DRAG DROP -

Drag and drop the SNMP attributes in Cisco IOS devices from the left onto the correct SNMPv2c or SNMPv3 categories on the right.

Select and Place:

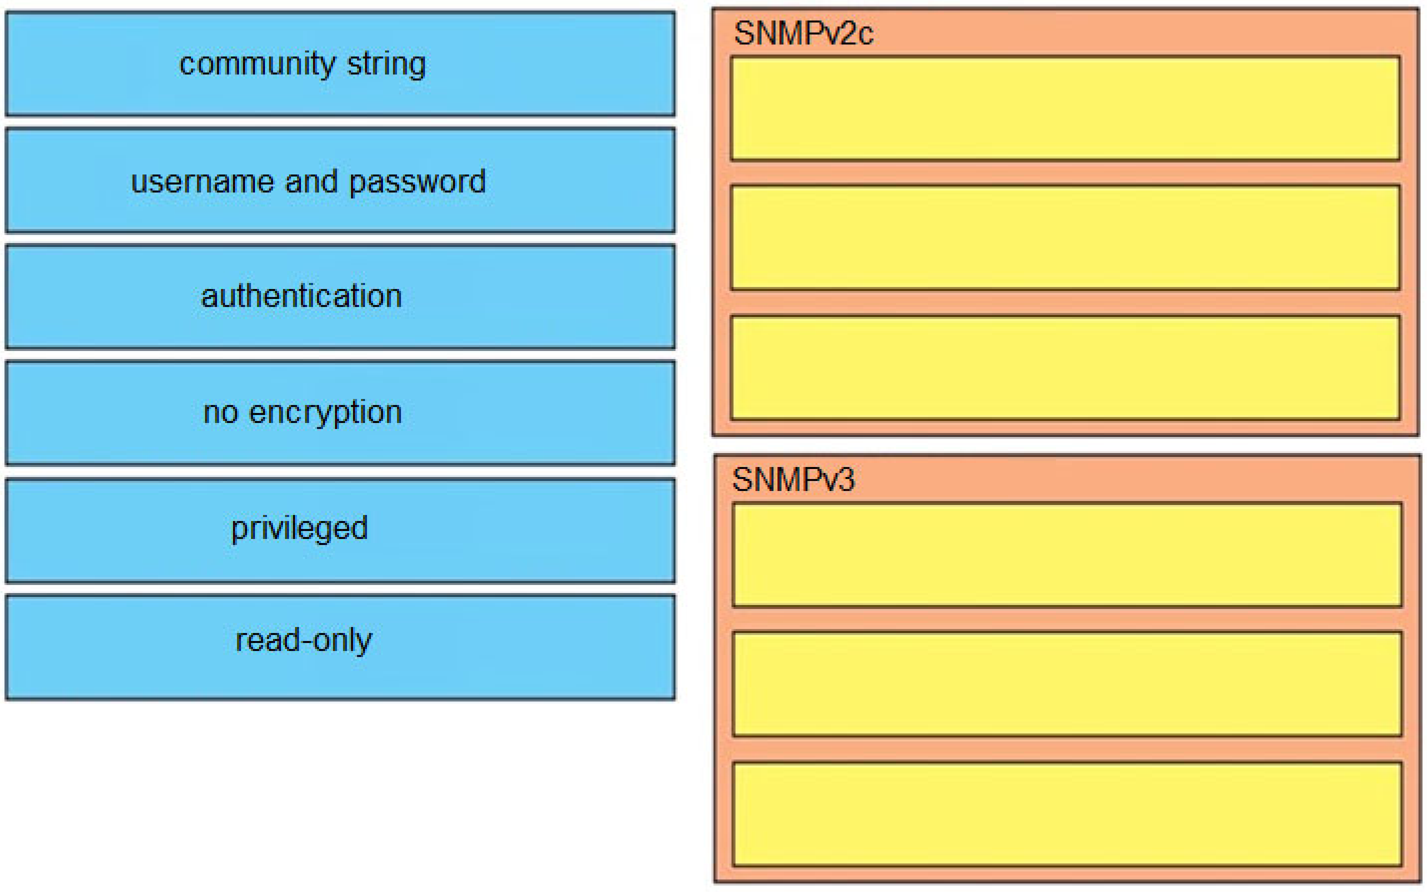

community string
username and password
authentication
no encryption
privileged
read-only

SNMPv2c

SNMPv3

Correct Answer:

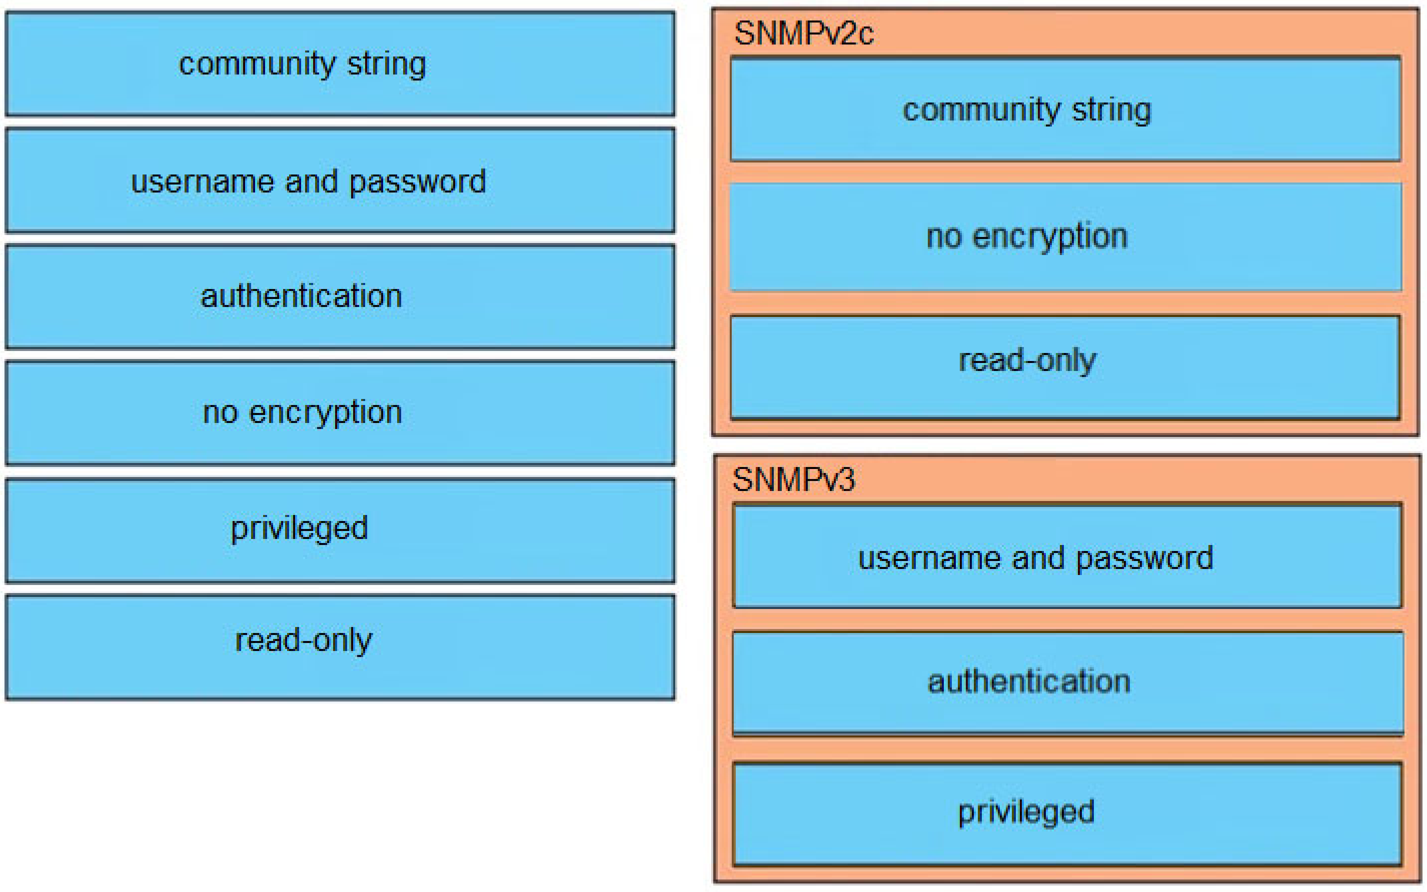

community string
username and password
authentication
no encryption
privileged
read-only

SNMPv2c
community string
no encryption
read-only

SNMPv3
username and password
authentication
privileged

⊟ **inteldarvid** 5 months, 2 weeks ago

Correct

upvoted 1 times

⊟ **HungarianDish** 7 months, 3 weeks ago

Imho, the given answer is correct. Users only come with SNMPv3, until that just community strings. SNMPv2 -> we can set read-only or read-write.

https://study-ccna.com/snmpv3-overview-configuration/
https://networklessons.com/cisco/ccnp-encor-350-401/how-to-configure-snmpv3-on-cisco-ios-router
https://www.examguides.com/Retired/ccna/200-125/cisco-ccna-50.htm
https://www.youtube.com/watch?v=hP5yA3hJlAc
https://www.youtube.com/watch?v=9Vx16VqzS8c

upvoted 2 times

⊟ **Zizu007** 11 months, 2 weeks ago

answer is correct. there is no such concept of user/pass in SNMPv2. Rather community-string.

upvoted 4 times

⊟ **xzckk** 1 year ago

The answer is wrong. It should be SNMPv2c -- username and password. SNMPv3 read-only.

upvoted 2 times

Question #176 Topic 1

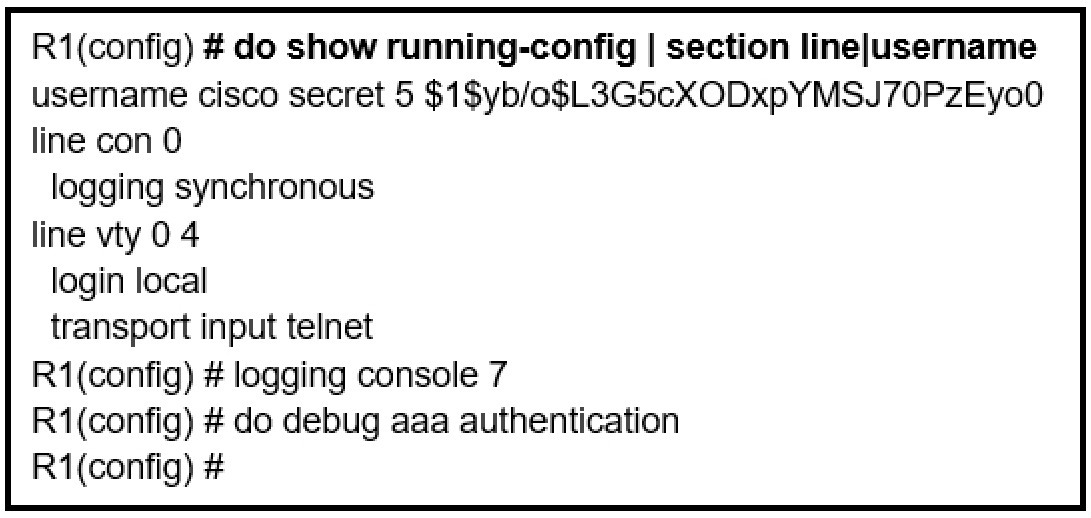

```
R1(config) # do show running-config | section line|username
username cisco secret 5 $1$yb/o$L3G5cXODxpYMSJ70PzEyo0
line con 0
  logging synchronous
line vty 0 4
  login local
  transport input telnet
R1(config) # logging console 7
R1(config) # do debug aaa authentication
R1(config) #
```

Refer to the exhibit. An administrator that is connected to the console does not see debug messages when remote users log in.
Which action ensures that debug messages are displayed for remote logins?

A. Enter the transport input ssh configuration command.

B. Enter the terminal monitor exec command.

C. Enter the logging console debugging configuration command.

D. Enter the aaa new-model configuration command.

Correct Answer: *C*

Community vote distribution

D (100%)

Alnet Highly Voted 2 years ago

Hold up. Let's look into this.
A isn't going to do anything for this problem.
B Term Mon is ONLY to display (shunt/pipe) console messages TO VTY LINES. Since we're connected to the Console, this will have no effect.
C Logging Console Debugging command was already entered... Logging Console 7 is the same command, you can use the severity level (0-7) or you can use the fancy name (debug, err, info...). They have the same effect. https://learningnetwork.cisco.com/s/question/0D53i00000Kt78L/no-logging-console-vs-no-logging-console-debug
D This IS a requirement to see AAA debug messages, there's nothing that indicates it's been entered yet, although the show run command would filter it out if it were already entered. But it's still the most likely candidate.
Answer is D.

upvoted 7 times

leecharxos 1 year, 11 months ago

also https://community.cisco.com/t5/networking-documents/how-to-configure-logging-in-cisco-ios/ta-p/3132434

upvoted 2 times

guy276465281819372 Most Recent 4 months, 2 weeks ago

All of the answers are not 100% correct BUT.
1. We do not know weather the users log in via SSH or Telnet
2. only for remote vty logging, the admin is using console port.
3. Already configured
4. only viable answer.

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

D correct

upvoted 1 times

HungarianDish 7 months, 3 weeks ago

I can't tell from the output whether using AAA is intended. If yes, then D is required. Maybe the question is formed differently on the real exam.

upvoted 1 times

PimplePooper 11 months, 2 weeks ago

Selected Answer: D

Answer C is already configured. Based on the remaining answers, D makes more sense.

upvoted 1 times

Nonono 1 year, 10 months ago

Selected Answer: D

C is already configured
upvoted 1 times

Nonono 1 year, 10 months ago
C is already configured. Answer D is correct.
upvoted 2 times

wts 1 year, 11 months ago
A - we are talking about "remote logins", including ssh, but only telnet is allowed.
B - does not make sense, since the administrator is connected via the console (cable).
C - what is needed for debug messages to be displayed when connected via the console.
D - why do you need to set it up, it does not affect anything.

Answer is C.
upvoted 3 times

OhBee 1 year, 11 months ago
C is already configured though...logging console 7 is the same as logging console debugging
upvoted 4 times

wts 1 year, 11 months ago
There should still be an answer.
upvoted 1 times

wts 1 year, 11 months ago
ANSWER IS D
Sorry. The debug aaa authentication command is run.
upvoted 1 times

wts 1 year, 10 months ago
It looks like you are right. By default, debug messages are already in the console lines.

The answer is based on the assumption that these default settings have not been changed and the aaa command has not been entered.
upvoted 1 times

myrmike 1 year, 11 months ago
Maybe someone can enlighten me. I labbed this and as is I see login/logout messages on the console when trying to remote into the router. Besides when someone has remoted into the router what is being looked for?
upvoted 1 times

Alex147 1 year, 11 months ago
Selected Answer: D
C is only for VTY connections.
D is correct - aaa need new-model need to be enabled in configuration.
upvoted 3 times

studybuddy10 2 years, 1 month ago
D. labbed to confirm - no messages until aaa new-model is added.
upvoted 3 times

Raider1 2 years, 2 months ago
The correct answer is B. Terminal monitor:This command enables the display of debugging messages and system error messages for the current terminal (i.e., VTY or asynchronous line) session.
upvoted 1 times

[Removed] 1 year, 10 months ago
Yea only when you remote in (SSH/Tel using VTY line) here it clearly states he is consoled in so that eliminates B...
upvoted 1 times

error_909 2 years, 3 months ago
The Correct Answer is B
The question is very clear and fouces on the remote connections:
To enable remote connections using Telnet or SSH we must only use "Terminal monitor" in the EXEC mode.

To enable it for the console line "Loggin console" in the global config mode.

Loggin Synchronisation is only needed to make the debug result not to mix with the command that you are write at the admin at the same time.
upvoted 2 times

error_909 2 years, 2 months ago
Sorry its D
upvoted 2 times

examShark 2 years, 4 months ago

D is the correct answer

A. Enter the transport input ssh configuration command.
>its telnet
B. Enter the terminal monitor exec command.
>console already monitors
C. Enter the logging console debugging configuration command.
>already done with logging console 7
D. Enter the aaa new-model configuration command.
>only one left - lab'd it also

upvoted 4 times

ITBiscuit 2 years, 8 months ago

The answer is D .. I labbed it. C is not correct because the logging console command was already used (logging console 7 - level 7 is debugging thus we would be applying the same command twice.)

upvoted 2 times

Sadist1111 2 years, 11 months ago

D is correct. When you are connected to device via console, then you already has configured "terminal monitor by default". You need to enable aaa services.

upvoted 2 times

Sheet 2 years, 11 months ago

D is correct.

upvoted 1 times

Question #177 Topic 1

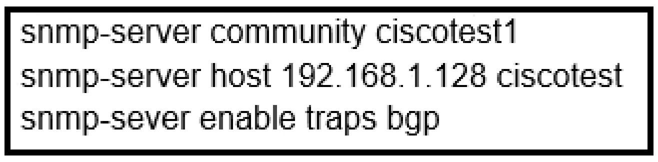

```
snmp-server community ciscotest1
snmp-server host 192.168.1.128 ciscotest
snmp-sever enable traps bgp
```

Refer to the exhibit. Network operations cannot read or write any configuration on the device with this configuration from the operations subnet.

Which two configurations fix the issue? (Choose two.)

A. Configure SNMP rw permission in addition to community ciscotest.

B. Modify access list 1 and allow operations subnet in the access list.

C. Modify access list 1 and allow SNMP in the access list.

D. Configure SNMP rw permission in addition to version 1.

E. Configure SNMP rw permission in addition to community ciscotest 1.

Correct Answer: *AB*

Community vote distribution

BE (73%) 13% 13%

anonymous1966 Highly Voted 3 years, 5 months ago

For me A and B is correct.
Setup SNMP Community with access-list
The best current practices recommend applying Access Control Lists (ACLs) to community strings and ensuring that the requests community strings are not identical to notifications community strings. Access lists provide further protection when used in combination with other protective measures.

This example sets up ACL to community string:

access-list 1 permit 1.1.1.1
snmp-server community string1 ro 1

Ref: https://www.cisco.com/c/en/us/support/docs/ip/simple-network-management-protocol-snmp/20370-snmpsecurity-20370.html

upvoted 10 times

DaanB Highly Voted 2 years, 8 months ago

The way the question is set, there are no correct answers. Based on the configuration, the communicty is ciscotest1. There is no access list 1 in the configuration. None of the answers follow this setup. If there would be an space between ciscotest and 1, then A and B would be correct - IMO

upvoted 9 times

MJM1973 Most Recent 3 weeks, 2 days ago

CORRECT ANSWER B and E
access-list 10 deny any
snmp-server host 10.1.1.1 mystring1
snmp-server community mystring1 RO 10

upvoted 1 times

Chiaretta 5 months ago

Selected Answer: AB

A and B are correct but missing the space on the question ciscotest 1, where 1 is the ACL number.

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: BE

B and E is correct, but only dependent if the question the syntax is ok:
snmp-server community cisco test (don't have acl)
or
snmp-server community cisco test 1 (different before, have acl)

upvoted 1 times

potato_inet0 7 months, 2 weeks ago

The question is written wrong.

It's snmp-server community ciscotest 1 where 1 is the ACL, otherwise the question does not make sense.

Taking the correction in account, the correct answer is A and B, because RO/RW are referenced before the ACL in the sintax, so answer E is wrong.

upvoted 1 times

Malasxd 7 months, 2 weeks ago

The question here is. The community name is ciscotest1 or it is ciscotest and the number 1 at the end is the ACL?

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

Based on this, password "ciscotest" + ACL 1.
https://www.cisco.com/c/en/us/support/docs/ip/simple-network-management-protocol-snmp/13506-snmp-traps.html#anc13

```
snmp-server community (community-string)
snmp-server host x.x.x.x (community-string, same as in snmp-server community)
```

In this case, the community-string = "ciscotest"

upvoted 1 times

HungarianDish 7 months, 3 weeks ago

Selected Answer: BE

If the output and E) contain community string ciscotest and access-list number 1, then BE is the closest for me. Probably they meant something like this:

```
#access-list 1 permit <operations subnet>
#snmp-server community ciscotest RW 1
```

upvoted 2 times

Dacusai 7 months, 3 weeks ago

B&E, the community name is ciscotest1, so E is the correct one, and add the subnet to the access list.

upvoted 2 times

forccnp 9 months ago

Selected Answer: B

B and E are correct answers

upvoted 2 times

forccnp 9 months, 1 week ago

Selected Answer: BE

B&E are correct answers

upvoted 1 times

Noproblem22 1 year ago

BE are the best answers

upvoted 2 times

tamangao 1 year, 1 month ago

Answer A is correct not E. you CANNOT specify the rw parameter after the ACL

```
R4(config)#snmp-server community ciscotest ?
<1-99> Std IP accesslist allowing access with this community string
<1300-1999> Expanded IP accesslist allowing access with this community
string
WORD Access-list name
ipv6 Specify IPv6 Named Access-List
ro Read-only access with this community string
rw Read-write access with this community string
view Restrict this community to a named MIB view
<cr> <cr>

R4(config)#snmp-server community ciscotest 1 ?
<cr> <cr>
```

upvoted 2 times

Huntkey 1 year, 2 months ago

Selected Answer: AB

There gotta be a space between ciscotest and 1, which makes "1" the ACL. If not, then E alone will be fine. Why asking to choose two answers?

upvoted 1 times

baid 1 year, 9 months ago

I think it is B E, A don't use the defined ACL, E use the defined ACL.

upvoted 1 times

wts 1 year, 9 months ago

Selected Answer: BE

A - it's about sending traps. The question describes a different problem.
B - just because you have to choose another option. It can be assumed (this is not in the question) that somewhere in the config there is an ACL. Well, in this case, if it is not allowed there, it would be necessary to allow snmp and the operations subnet. B because the question indicates that the subnet is operations.
C - well, I chose B...

D - nothing indicates that it is needed.
E - if not explicitly specified, then at the end of the snmp-server community command there will be an implicit RO. This needs to be fixed.

upvoted 2 times

⊟ **timtgh** 1 year, 6 months ago

A and E seem identical, except E has a 1 at the end.

upvoted 1 times

⊟ **wts** 1 year, 3 months ago

If we choose A, then we don't need anything else.

And the author of the question asks to choose two options.

upvoted 1 times

⊟ **Hack4** 1 year, 10 months ago

BE are the best answer

upvoted 1 times

Question #178

Topic 1

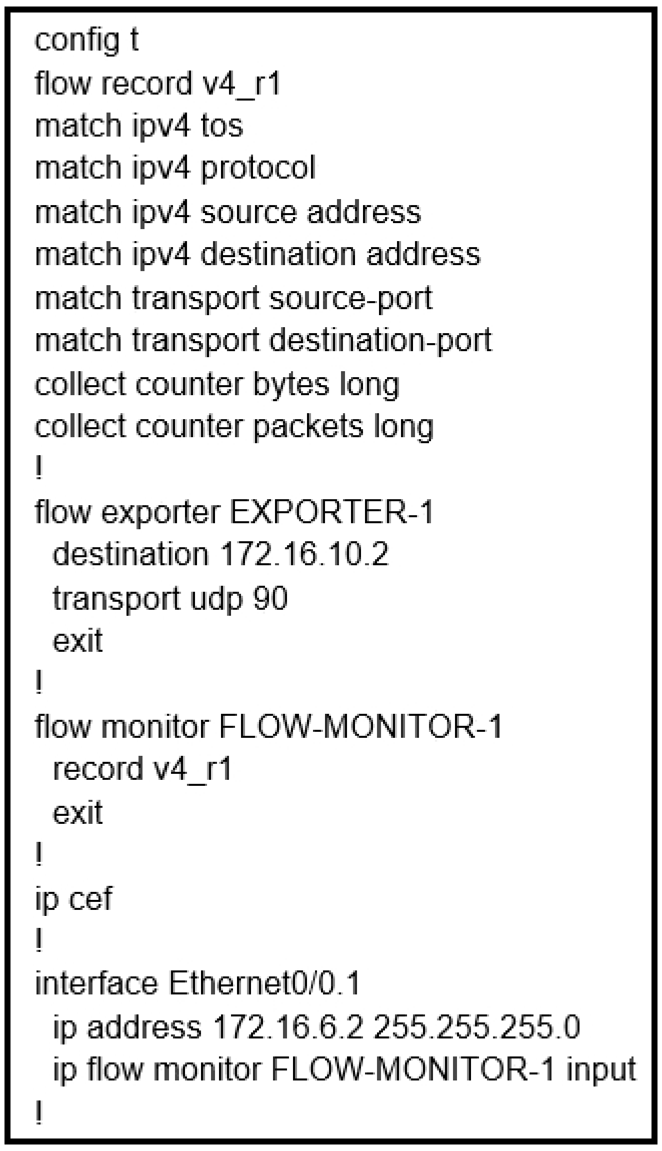

```
config t
flow record v4_r1
match ipv4 tos
match ipv4 protocol
match ipv4 source address
match ipv4 destination address
match transport source-port
match transport destination-port
collect counter bytes long
collect counter packets long
!
flow exporter EXPORTER-1
 destination 172.16.10.2
 transport udp 90
 exit
!
flow monitor FLOW-MONITOR-1
 record v4_r1
 exit
!
ip cef
!
interface Ethernet0/0.1
 ip address 172.16.6.2 255.255.255.0
 ip flow monitor FLOW-MONITOR-1 input
!
```

Refer to the exhibit. Why is the remote NetFlow server failing to receive the NetFlow data?

A. The flow exporter is configured but is not used.

B. The flow monitor is applied in the wrong direction.

C. The flow monitor is applied to the wrong interface.

D. The destination of the flow exporter is not reachable.

Correct Answer: *D*

Community vote distribution

A (100%)

S_E_T Highly Voted 3 years, 6 months ago

The correct answer is A.
The exporter is not configured under the flow monitor.
https://www.cisco.com/c/dam/en/us/td/docs/security/stealthwatch/netflow/Cisco_NetFlow_Configuration.pdf

upvoted 21 times

CraigB83 Highly Voted 3 years, 2 months ago

A is correct

flow exporter EXPORTER-1
destination 172.16.10.2
transport udp 90
exit
flow monitor FLOW-MONITOR-1
record netflow ipv4 original-input
exporter EXPORTER-2
exporter EXPORTER-1

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/fnetflow/configuration/15-mt/fnf-15-mt-book/cfg-de-fnflow-exprts.html

upvoted 11 times

inteldarvid Most Recent 5 months, 2 weeks ago

Selected Answer: A

Correct A

upvoted 1 times

Calyfas 9 months, 3 weeks ago

Selected Answer: A

A is correct

upvoted 1 times

ERICKPORRAS 1 year, 3 months ago

Selected Answer: A

A is correct

upvoted 1 times

phryde 1 year, 6 months ago

Selected Answer: A

A is correct

upvoted 1 times

Bruffas 1 year, 9 months ago

Selected Answer: A

Has to be A

upvoted 1 times

Nonono 1 year, 10 months ago

Selected Answer: A

tested A

upvoted 1 times

Jenia1 1 year, 11 months ago

Selected Answer: A

A is correct

upvoted 1 times

Kai12345 2 years ago

Selected Answer: A

A is correct

upvoted 1 times

studybuddy10 2 years, 1 month ago

A correct.
Simple steps needed:
1. Create record
2. Create Exporter
3 Create monitor and reference record and exporter
4 assign monitor to an interface

upvoted 1 times

OakA1 2 years, 2 months ago

Only A can be correct

upvoted 1 times

error_909 2 years, 3 months ago

The correct answer is A.
The flow exporter is configured but is not used

upvoted 1 times

examShark 2 years, 4 months ago

The correct answer is A

upvoted 1 times

RemiK 2 years, 5 months ago

A is definitely the correct answer.
Thanks to S_E_T for the link that confirms it.

upvoted 1 times

ssbipa6 2 years, 8 months ago

A is correct indeed

upvoted 1 times

Sheet 2 years, 11 months ago

A is correct

upvoted 1 times

Question #179

Topic 1

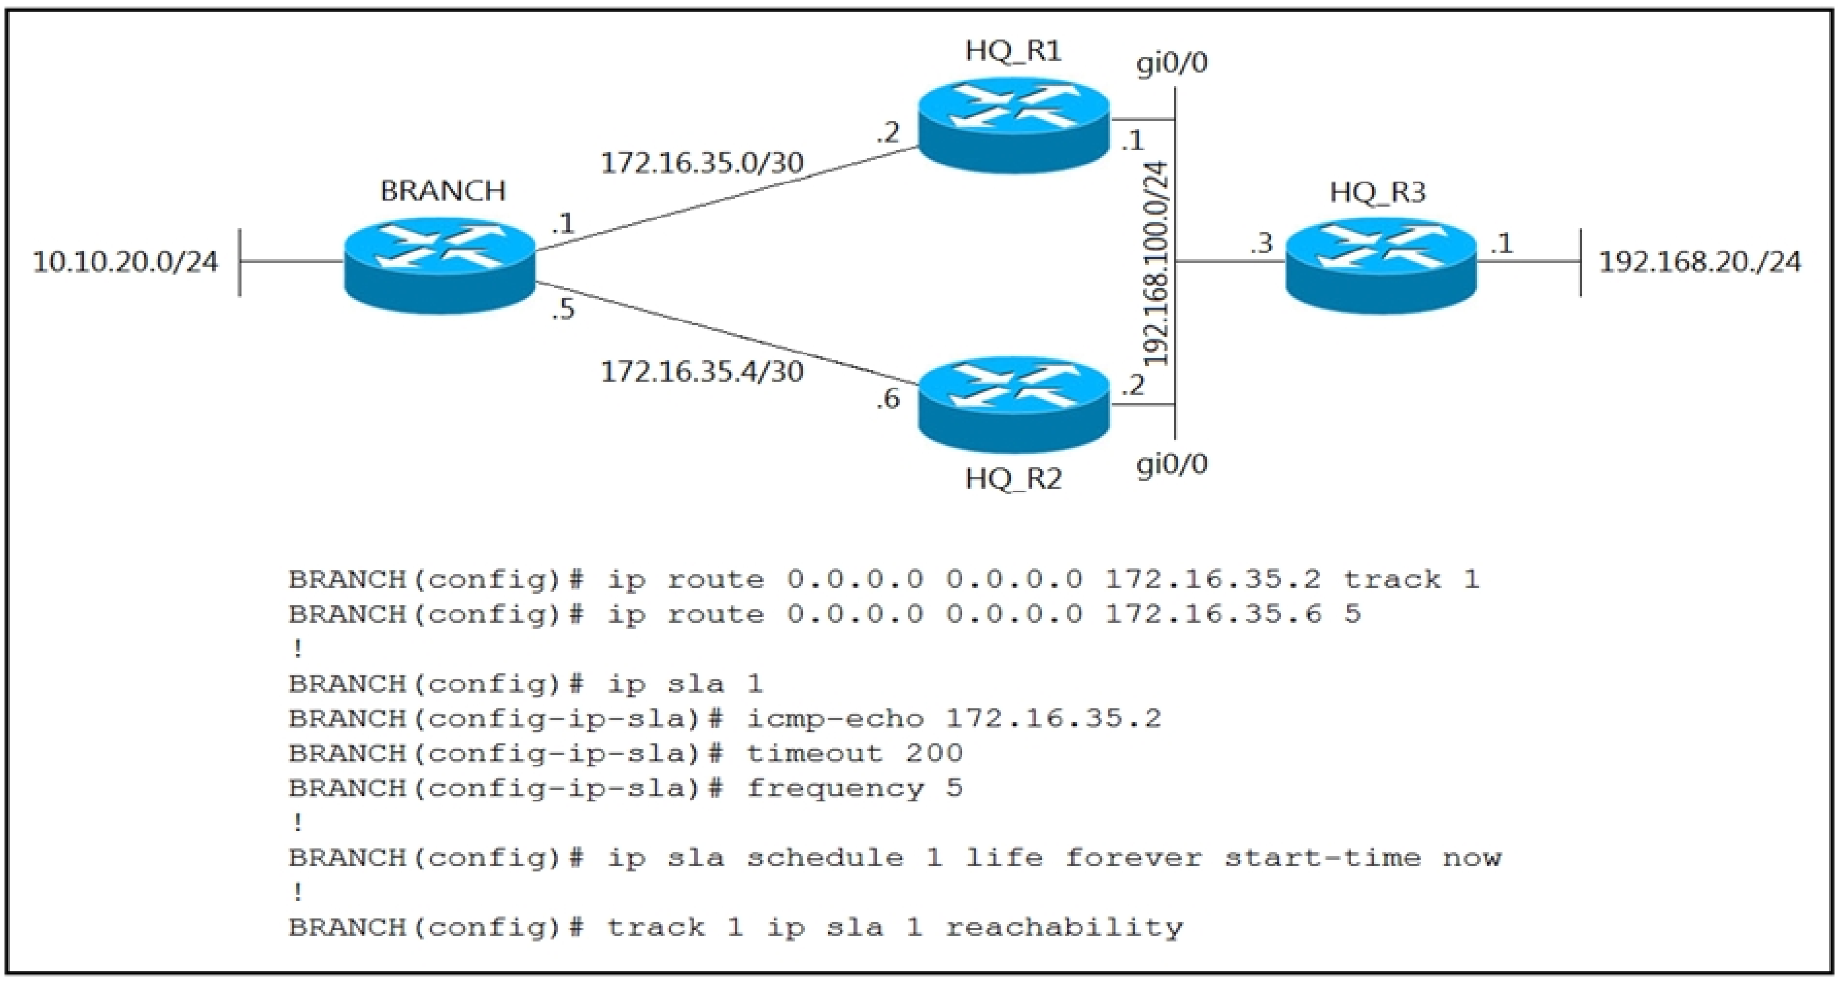

```
BRANCH(config)# ip route 0.0.0.0 0.0.0.0 172.16.35.2 track 1
BRANCH(config)# ip route 0.0.0.0 0.0.0.0 172.16.35.6 5
!
BRANCH(config)# ip sla 1
BRANCH(config-ip-sla)# icmp-echo 172.16.35.2
BRANCH(config-ip-sla)# timeout 200
BRANCH(config-ip-sla)# frequency 5
!
BRANCH(config)# ip sla schedule 1 life forever start-time now
!
BRANCH(config)# track 1 ip sla 1 reachability
```

Refer to the exhibit. An engineer has successfully set up a floating static route from the BRANCH router to the HQ network using HQ_R1 as the primary default gateway. When the g0/0 goes down on HQ_R1, the branch network cannot reach the HQ network 192.168.20.0/24.
Which configuration resolves the issue?

A. HQ_R3(config)# ip sla responder HQ_R3(config)# ip sla responder icmp-echo 172.16.35.1

B. BRANCH(config)# ip sla 1 BRANCH(config-ip-sla)# icmp-echo 192.168.100.2

C. HQ_R3(config)# ip sla responder HQ_R3(config)# ip sla responder icmp-echo 172.16.35.5

D. BRANCH(config)# ip sla 1 BRANCH(config-ip-sla)# icmp-echo 192.168.100.1

Correct Answer: *D*

Community vote distribution

D (100%)

AlexInShort12 3 days, 22 hours ago

ip sla responder icmp-echo doesn't seems to be a real command.
Not what difference it make direct route vs 192.168.100.1

upvoted 1 times

Calyfas 9 months, 3 weeks ago

Selected Answer: D

D is correct, we need to monitor G0/0 from HQ_R1

upvoted 4 times

Question #180 Topic 1

An engineer configured a DHCP server for Cisco IP phones to download its configuration from a TFTP server, but the IP phones failed to load the configuration.
What must be configured to resolve the issue?

A. BOOTP port 67

B. DHCP option 66

C. BOOTP port 68

D. DHCP option 69

Correct Answer: *B*

Community vote distribution

B (100%)

glbngl91 Highly Voted 11 months, 1 week ago
"Commander DHCP, the IPPhone has come... Execute Option 66"
Correct answer, btw
upvoted 5 times

inteldarvid Most Recent 5 months, 2 weeks ago
Selected Answer: B
B is correct. :

DHCP option 150 provides the IP addresses of a list of TFTP servers.
DHCP option 66 gives the IP address or the hostname of a single TFTP server.

https://www.cisco.com/c/en/us/td/docs/security/asa/asa84/configuration/guide/asa_84_cli_config/basic_dhcp.html
upvoted 1 times

JingleJangus 1 year, 7 months ago
Most pages online seem to suggest something similar to the following:
Option 66: Provides the address of a single TFTP server
Option 150: Provides a list of multiple TFTP server addresses
upvoted 2 times

Bruffas 1 year, 9 months ago
Selected Answer: B
The given answer is correct
upvoted 1 times

Mjestic 2 years, 3 months ago
This question seems weird. I have never saw option 66 for Cisco IP Phones, always option 150.
And I just checked on different blogs, option 66 is mostly for Juniper and option 150 is for Cisco. But it seems that for very rare cases, option 66 can be used when option 150 is not available. Sad question again...
upvoted 3 times

Stupid 1 year, 11 months ago
I agree, the best link I could find about it is page 6 on
https://www.cisco.com/c/en/us/td/docs/security/asa/asa84/configuration/guide/asa_84_cli_config/basic_dhcp.pdf
upvoted 2 times

jarz 1 year, 3 months ago
Here's another good explication. I fucking hate these certification exams with these types of bullshit questions.

https://blog.router-switch.com/2013/03/dhcp-option-150-dhcp-option-66/
upvoted 5 times

rggod 2 years ago
I had this in my notes, Opt 66 uses TFTP w/ hostnames. Opt 150 is same but uses IP addresses which is why it's more common to see.
upvoted 3 times

examShark 2 years, 4 months ago
Thye given answer is correct
upvoted 1 times

Question #181 Topic 1

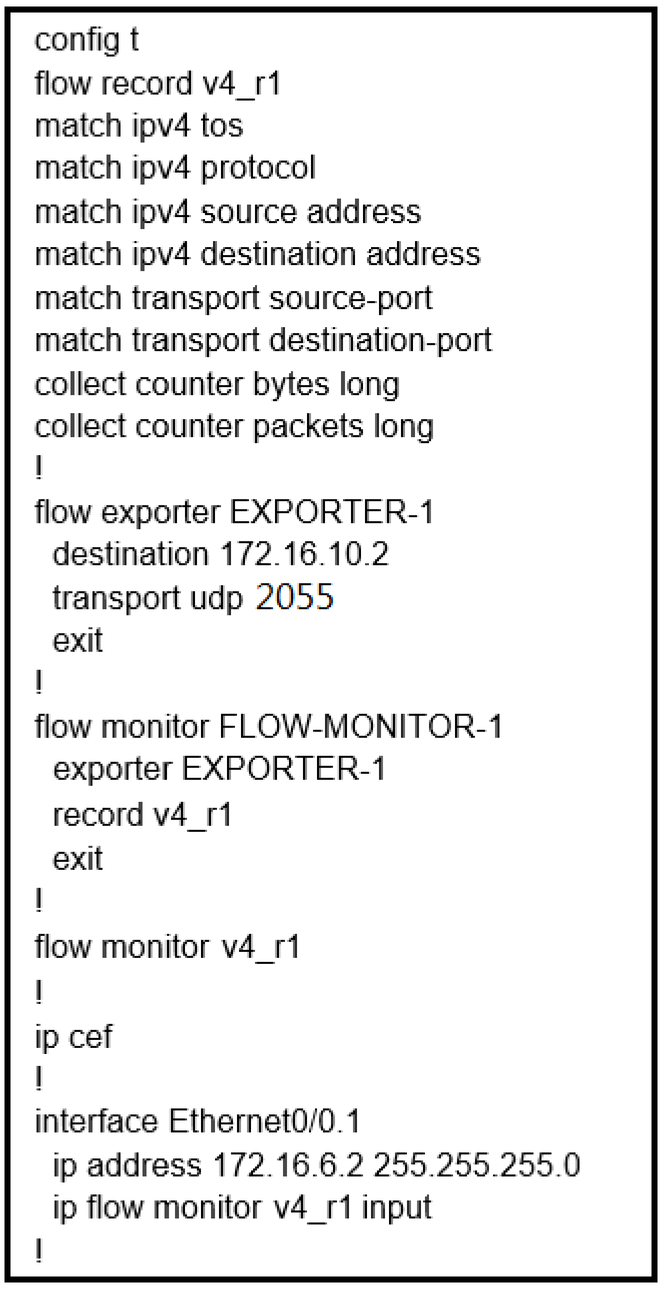

```
config t
flow record v4_r1
match ipv4 tos
match ipv4 protocol
match ipv4 source address
match ipv4 destination address
match transport source-port
match transport destination-port
collect counter bytes long
collect counter packets long
!
flow exporter EXPORTER-1
  destination 172.16.10.2
  transport udp 2055
  exit
!
flow monitor FLOW-MONITOR-1
  exporter EXPORTER-1
  record v4_r1
  exit
!
flow monitor v4_r1
!
ip cef
!
interface Ethernet0/0.1
  ip address 172.16.6.2 255.255.255.0
  ip flow monitor v4_r1 input
!
```

Refer to the exhibit. The remote server is failing to receive the NetFlow data.
Which action resolves the issue?

A. Modify the flow transport command transport udp 2055 to move under flow monitor profile.

B. Modify the interface command to ip flow monitor FLOW-MONITOR-1 input.

C. Modify the udp port under flow exporter profile to ip transport udp 4739.

D. Modify the flow record command record v4_r1 to move under flow exporter profile.

Correct Answer: *B*

Community vote distribution

B (100%)

Bruffas 1 year, 9 months ago

Selected Answer: B

The given answer is correct

upvoted 1 times

studybuddy10 2 years, 1 month ago

B - correct, FLOW-MONITOR-1 has a record and an exporter, v4_r1 has none.

upvoted 4 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #182 Topic 1

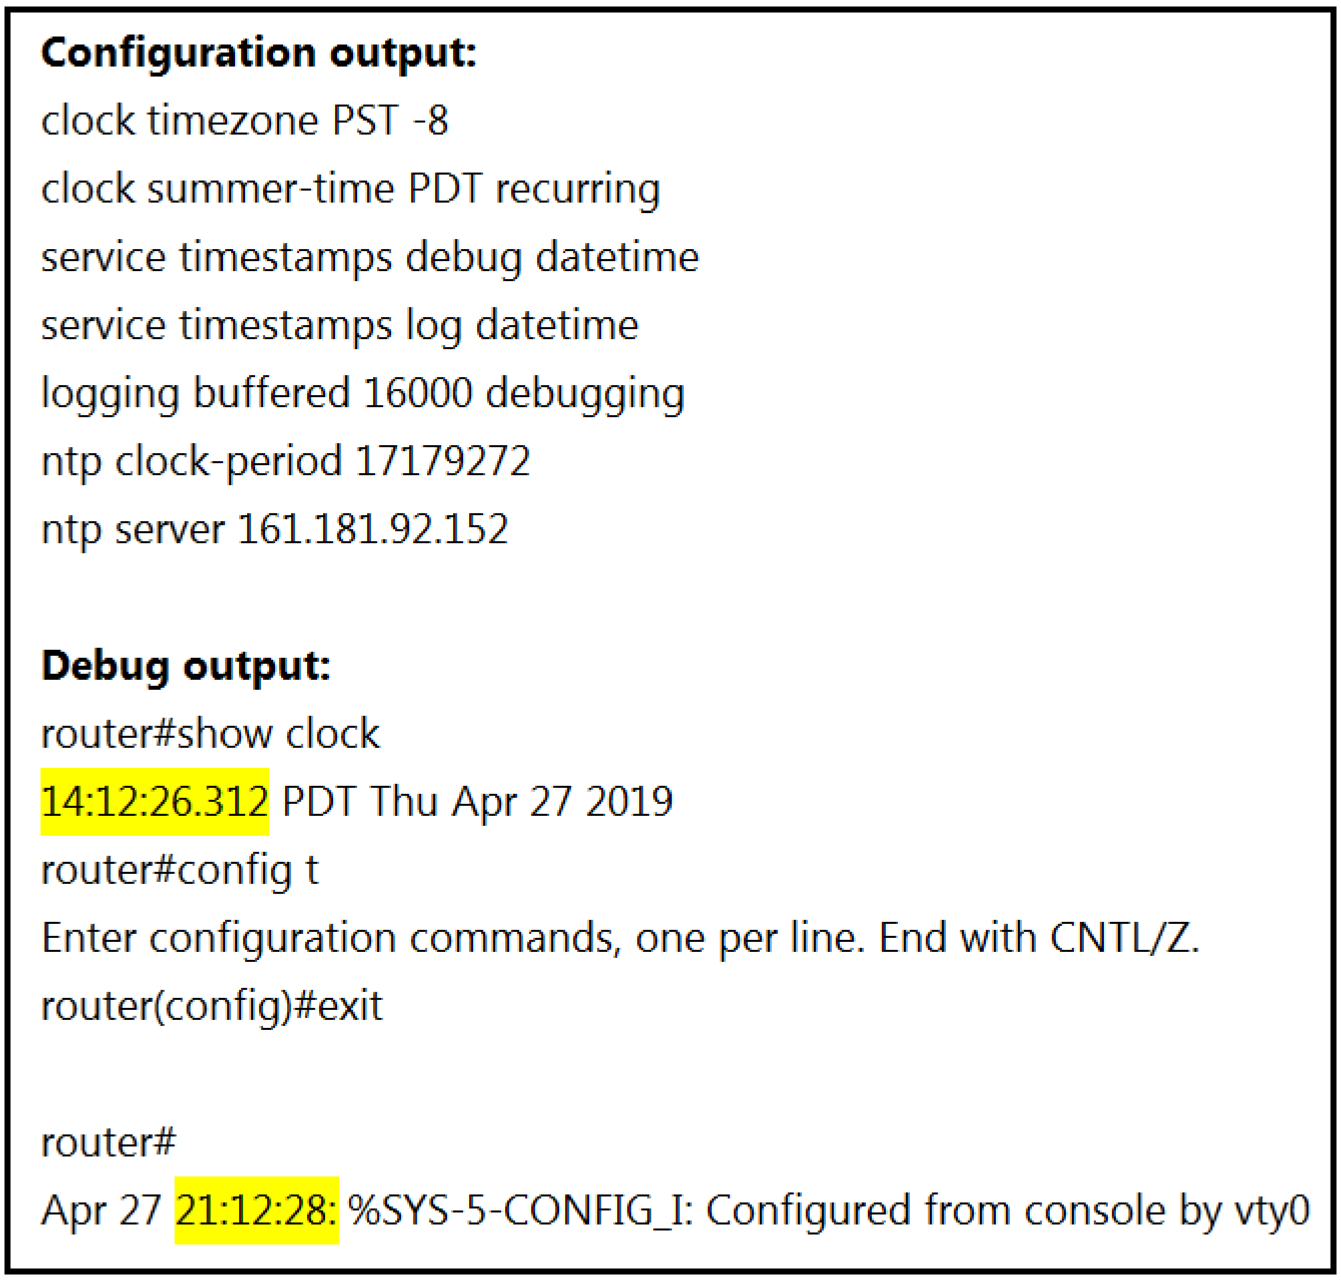

Configuration output:

```
clock timezone PST -8
clock summer-time PDT recurring
service timestamps debug datetime
service timestamps log datetime
logging buffered 16000 debugging
ntp clock-period 17179272
ntp server 161.181.92.152
```

Debug output:

```
router#show clock
14:12:26.312 PDT Thu Apr 27 2019
router#config t
Enter configuration commands, one per line. End with CNTL/Z.
router(config)#exit

router#
Apr 27 21:12:28: %SYS-5-CONFIG_I: Configured from console by vty0
```

Refer to the exhibit. A network administrator configured NTP on a Cisco router to get synchronized time for system and logs from a unified time source. The configuration did not work as desired.

Which service must be enabled to resolve the issue?

A. Enter the service timestamps log datetime clock-period global command.

B. Enter the service timestamps log datetime synchronize global command.

C. Enter the service timestamps log datetime console global command.

D. Enter the service timestamps log datetime localtime global command.

Correct Answer: *D*

Community vote distribution

D (100%)

Bruffas 1 year, 9 months ago

Selected Answer: D

The given answer is correct

upvoted 2 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #183 Topic 1

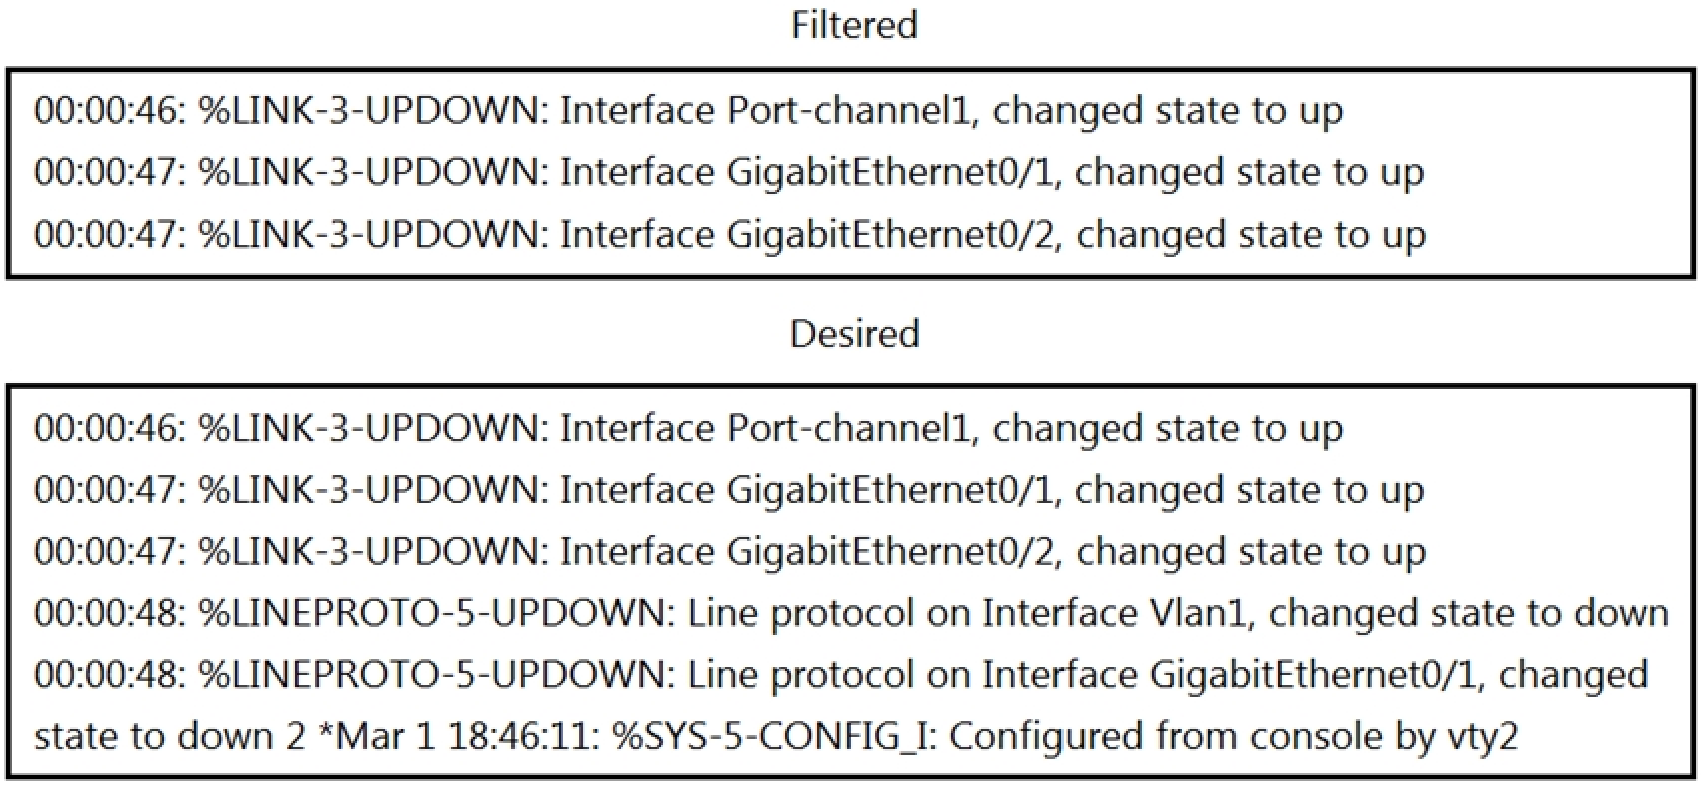

Filtered

```
00:00:46: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/1, changed state to up
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/2, changed state to up
```

Desired

```
00:00:46: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/1, changed state to up
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/2, changed state to up
00:00:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1, changed state to down
00:00:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/1, changed
state to down 2 *Mar 1 18:46:11: %SYS-5-CONFIG_I: Configured from console by vty2
```

Refer to the exhibits. An engineer filtered messages based on severity to minimize log messages. After applying the filter, the engineer noticed that it filtered required messages as well.

Which action must the engineer take to resolve the issue?

A. Configure syslog level 2.

B. Configure syslog level 3.

C. Configure syslog level 4.

D. Configure syslog level 5.

Correct Answer: *D*

Community vote distribution

D (100%)

HungarianDish 6 months, 3 weeks ago

Selected Answer: D

Specifying a level causes messages at that level and numerically lower levels to be displayed at the destination.
https://www.cisco.com/c/en/us/td/docs/routers/access/wireless/software/guide/SysMsgLogging.html

#logging trap 5
= send notifications and lower severity levels (5,4,3,2,1)

https://www.ciscopress.com/articles/article.asp?p=426638&seqNum=3

upvoted 1 times

Dacusai 7 months, 2 weeks ago

B is the correct one. When you filter something it means you eliminate it, so if you filter level 3 like the picture says and on the desired part are included, it means that you need to add Level 3 to the config no Level 5 because Level 5 is already there.

upvoted 1 times

marc2109 1 year, 1 month ago

The syslog level is given in the Desired logging exhibit: "%LINEPROTO-5-UPDOWN". So it's level 5.

upvoted 3 times

Bruffas 1 year, 9 months ago

Selected Answer: D

The given answer is correct
D

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

Question #184 *Topic 1*

An engineer is troubleshooting on the console session of a router and turns on multiple debug commands. The console screen is filled with scrolling debug messages that none of the commands can be verified if entered correctly or display any output.
Which action allows the engineer to see entered console commands while still continuing the analysis of the debug messages?

A. Configure the term no mon command globally.

B. Configure the logging synchronous level all command.

C. Configure the logging synchronous command.

D. Configure the no logging console debugging command globally.

Correct Answer: *C*

Community vote distribution

C (100%)

AliMo123 Highly Voted 2 years, 1 month ago

C is correct
"What is logging synchronous command in Cisco?
This command controls the printing of log messages to a user's terminal. By default, messages are printed at any time, possibly disrupting the user's current command. This command tells the router to wait until the user's current command and its output are completed before displaying any logging messages."

upvoted 7 times

inteldarvid Most Recent 5 months, 2 weeks ago

Selected Answer: C

C correct

upvoted 1 times

Bruffas 1 year, 9 months ago

Selected Answer: C

C is correct

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct
. Configure the logging synchronous command.

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

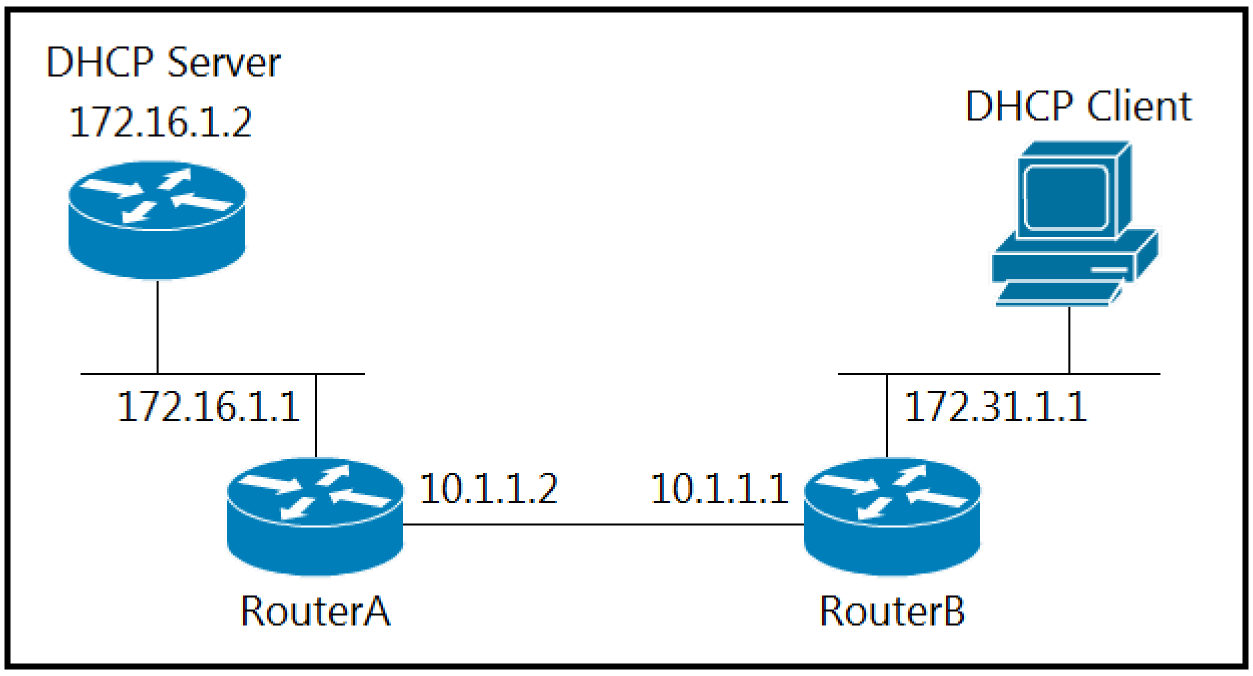

Refer to the exhibit. The DHCP client is unable to receive an IP address from the DHCP server. RouterB is configured as follows:

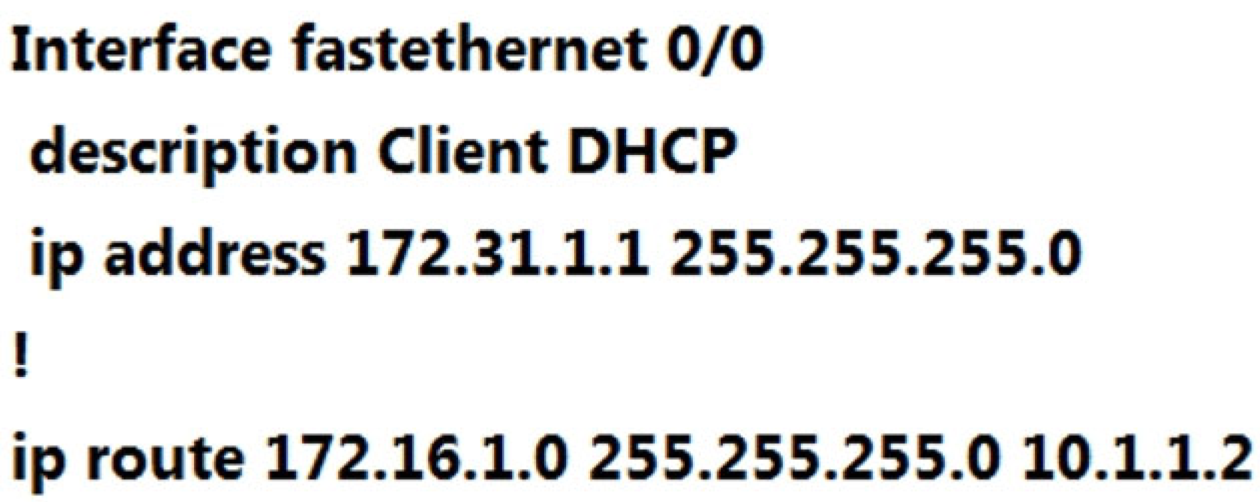

```
Interface fastethernet 0/0
 description Client DHCP
 ip address 172.31.1.1 255.255.255.0
!
ip route 172.16.1.0 255.255.255.0 10.1.1.2
```

Which command is required on the fastethernet 0/0 interface of RouterB to resolve this issue?

A. RouterB(config-if)#ip helper-address 172.16.1.1

B. RouterB(config-if)#ip helper-address 255.255.255.255

C. RouterB(config-if)#ip helper-address 172.16.1.2

D. RouterB(config-if)#ip helper-address 172.31.1.1

Correct Answer: *C*

Community vote distribution

C (100%)

⊟ jsanc1974 Highly Voted 2 years, 2 months ago

Answer C is correct because when using the ip helper-address command inside interface mode you have to add the ip address of the dhcp server that you are attempting to obtain dhcp information from

upvoted 5 times

⊟ inteldarvid Most Recent 5 months, 2 weeks ago

Selected Answer: C

C correct

upvoted 1 times

Question #186 Topic 1

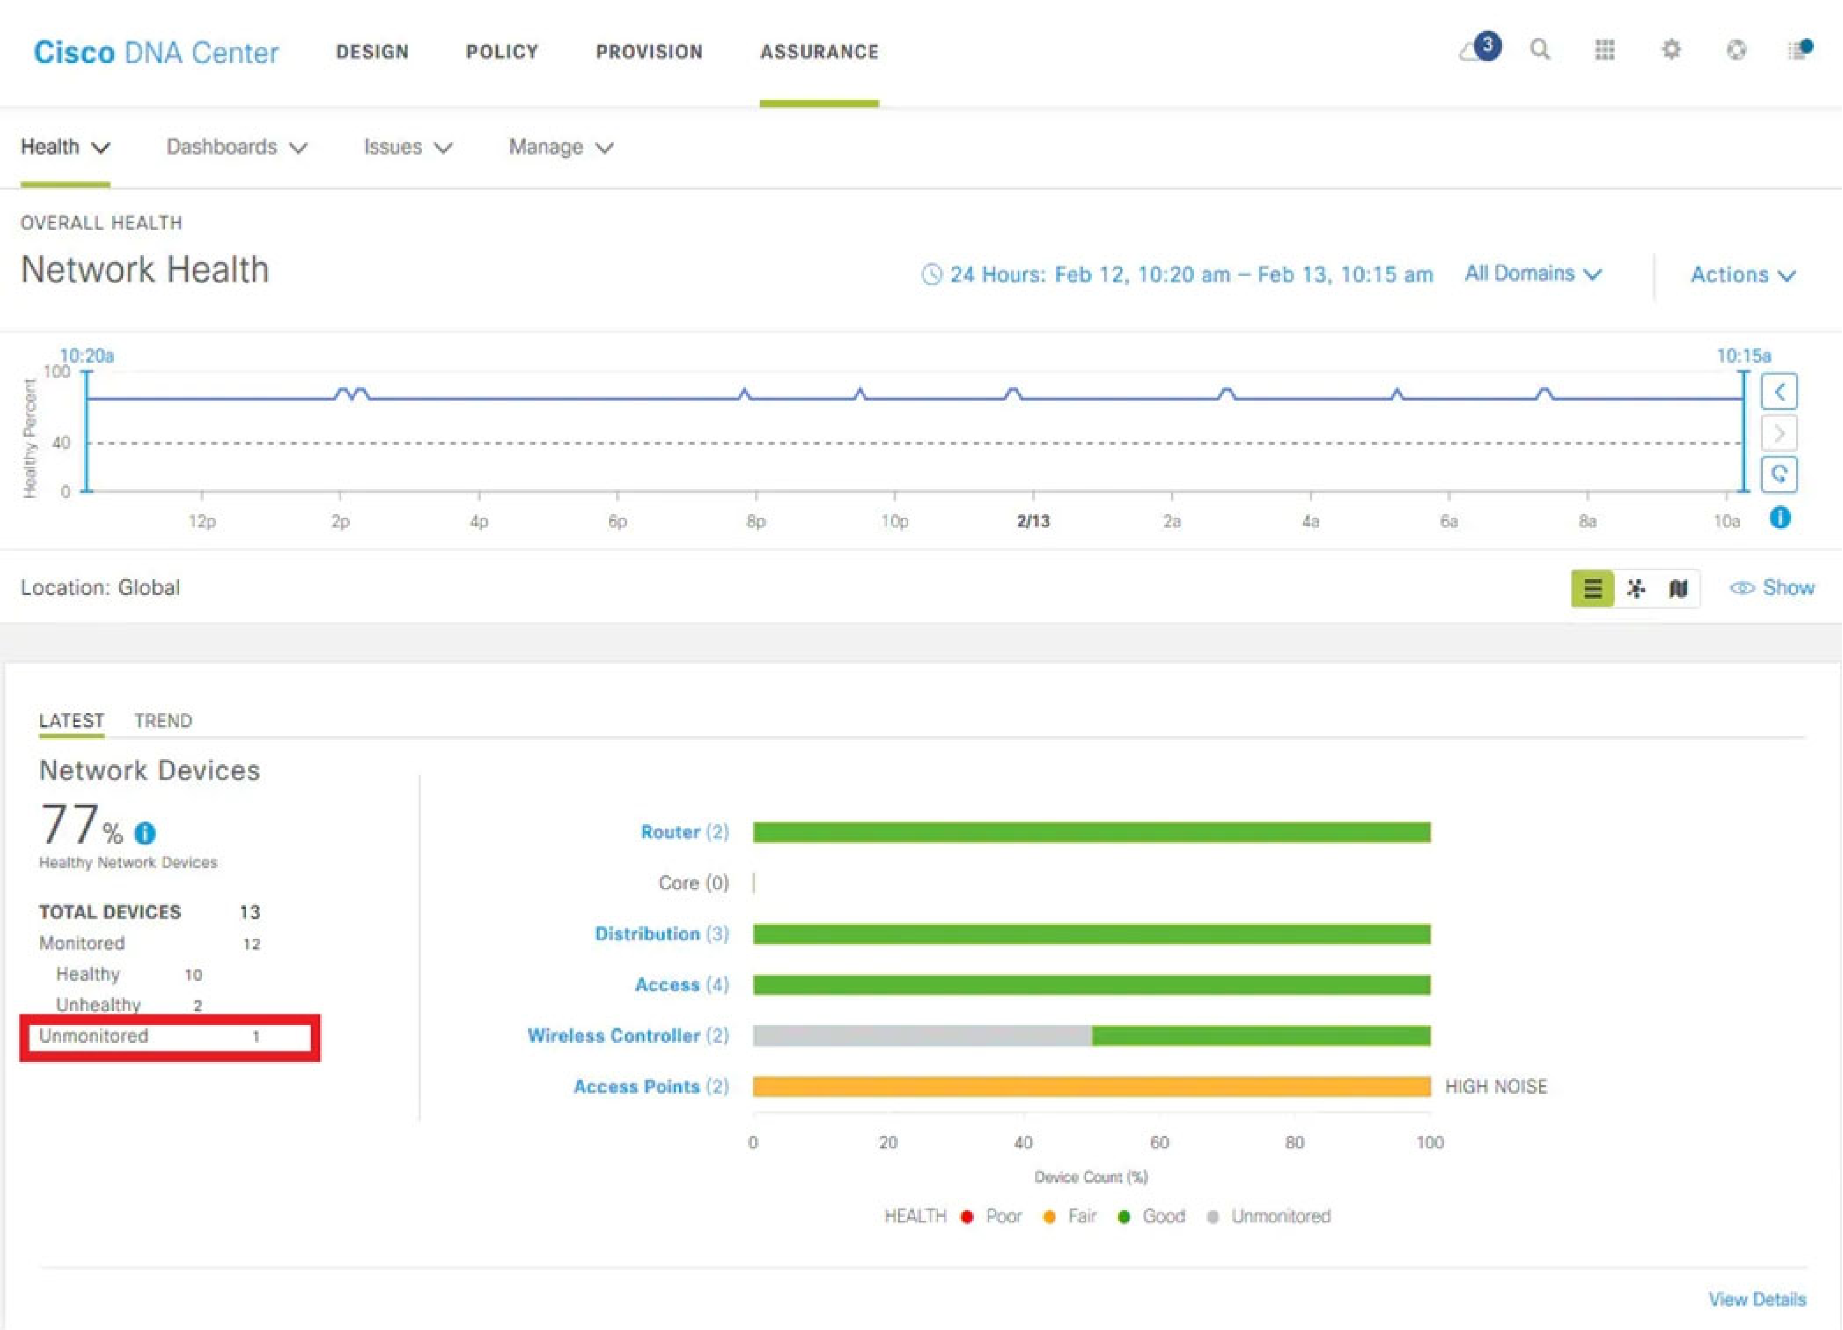

TOTAL DEVICES	13
Monitored	12
Healthy	10
Unhealthy	2
Unmonitored	1

Refer to the exhibit. A network administrator added one router in the Cisco DNA Center and checked its discovery and health from the Network Health Dashboard.

The network administrator observed that the router is still showing up as unmonitored.

What must be configured on the router to mount it in the Cisco DNA Center?

A. Configure router with SNMPv2c or SNMPv3 traps

B. Configure router with the telemetry data

C. Configure router with routing to reach Cisco DNA Center

D. Configure router with NetFlow data

Correct Answer: *B*

bjromero28 Highly Voted 2 years, 1 month ago

Unmonitored devices are devices for which Assurance did not receive any telemetry data during the specified time range. Unmonitored devices are included in the Network Health Score computation. They are used as part of the total number of devices against which the health device percentage is calculated.

Given Answer (B) is correct.

upvoted 7 times

Pietjeplukgeluk 2 weeks, 6 days ago

I understand you saying "did not receive any telemetry data". However, this is not the question. The question is how to resolve the issue with the addition of configuration. Can anyone explain me why B is correct, why should i configure a router with telemetry data? (i personally think the telemetry is generated at the router side, you will never CONFIGURE telemetry data, this is just a gathering of many data pointers over time).

upvoted 1 times

AlexInShort12 Most Recent 3 days, 22 hours ago

Multiple option could say this message.
To onboard a device, dna need SSH cred + SNMP and good routing...
To telemetry count has cred+ snmp.

upvoted 1 times

AlexInShort12 3 days, 22 hours ago

So probably B... since ssh cred is not there..

upvoted 1 times

Calyfas 9 months ago

Given Answer (B) is correct.

upvoted 1 times

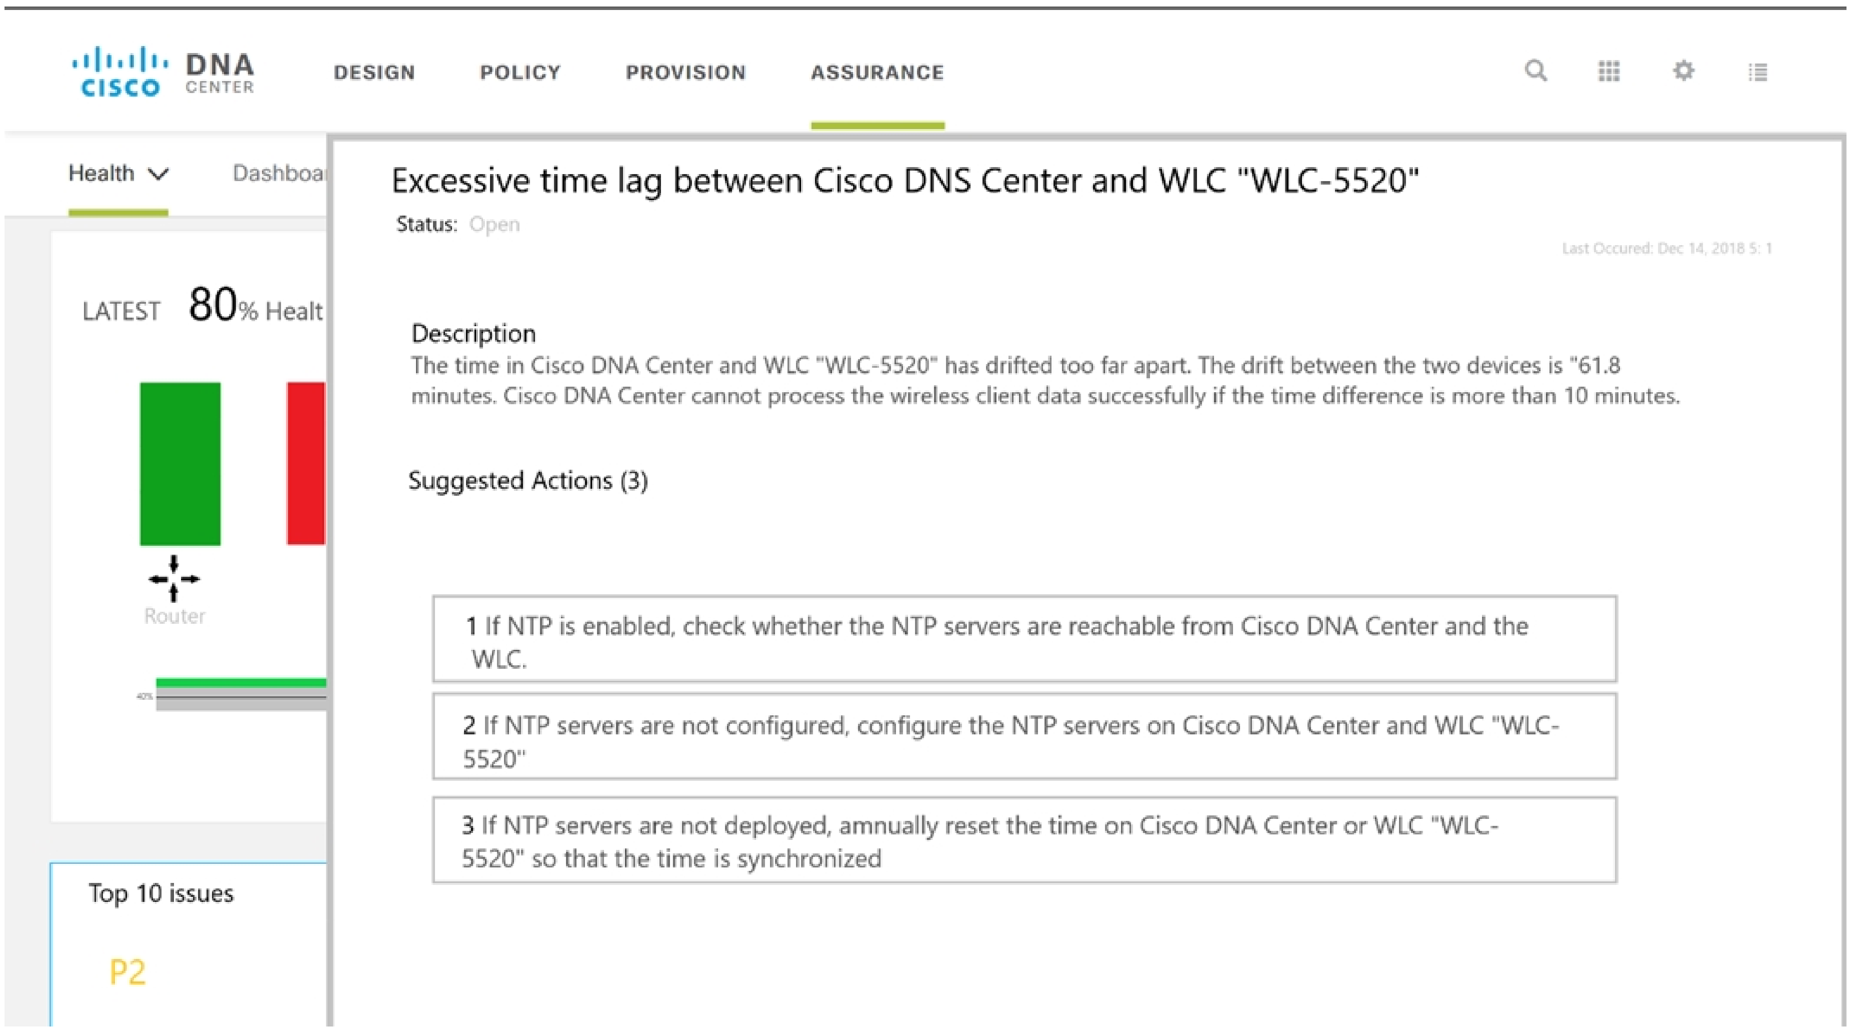

Excessive time lag between Cisco DNS Center and WLC "WLC-5520"

Status: Open

Last Occured: Dec 14, 2018 5: 1

Description

The time in Cisco DNA Center and WLC "WLC-5520" has drifted too far apart. The drift between the two devices is "61.8 minutes. Cisco DNA Center cannot process the wireless client data successfully if the time difference is more than 10 minutes.

Suggested Actions (3)

1 If NTP is enabled, check whether the NTP servers are reachable from Cisco DNA Center and the WLC.

2 If NTP servers are not configured, configure the NTP servers on Cisco DNA Center and WLC "WLC-5520"

3 If NTP servers are not deployed, annually reset the time on Cisco DNA Center or WLC "WLC-5520" so that the time is synchronized

Refer to the exhibit. NTP is configured across the network infrastructure and Cisco DNA Center. An NTP issue was reported on the Cisco DNA Center at 17:15.

Which action resolves the issue?

A. Reset the NTP server to resolve any synchronization issues for all devices

B. Check and resolve reachability between Cisco DNA Center and the NTP server

C. Check and resolve reachability between the WLC and the NTP server

D. Check and configure NTP on the WLC and synchronize with Cisco DNA Center

Correct Answer: *D*

Community vote distribution

C (70%) D (30%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: C

option C correct

upvoted 1 times

HungarianDish 7 months, 3 weeks ago

It seems to be a bug in DNA Center:
https://community.cisco.com/t5/cisco-digital-network-architecture-dna/dna-assurance-dna-center-and-network-device-time-has-drifted/td-p/4067331
https://bst.cloudapps.cisco.com/bugsearch/bug/CSCvr46035
"Workaround: Remove the timezone configuration from the network device. "

If it is about the mentioned bug then none of the answers fit. :(

upvoted 1 times

ntdevera 1 year, 3 months ago

Selected Answer: C

C
Not B, if DNA center cant reach NTP. It wont just be 1 WLC that has an issue.
Not D, You don't synchronize with the DNA Center to fix time issues. You sync it with the NTP server. It

upvoted 1 times

[Removed] 1 year, 4 months ago

Selected Answer: C

If NTP is already enabled across the infrastructure and only the WLC is reporting issues then its probably just an issue with the NTP server syncing with the WLC

upvoted 1 times

wts 1 year, 9 months ago

Selected Answer: D

I did not see a sufficient explanation here, but I myself can not give it.

Without thinking too much, I would exclude the answer options regarding reachability on the network. The problem is only in the time lag and only with the WLC. Therefore D.

...the JOKERR link is not working.

upvoted 3 times

bogd 1 year, 10 months ago

Selected Answer: C

With NTP already configured and only one device exhibiting the issue, I would first check that device and make sure that it can reach the NTP server and sync time.

upvoted 4 times

Dirkd0344 2 years ago

The question state that NTP is already configured on the network infrastructure. Next you would want to verify reachability to the NTP server from DNAC and the WLC. Therefore either A or C could be correct, but I would say A is the answer.

upvoted 3 times

[Removed] 1 year, 11 months ago

Eh, the first suggested action says to check for reachability so I would go with C since its the WLC having issues with communication..

upvoted 1 times

Jenia1 1 year, 10 months ago

Same for me, C seems to be the best option

A,B,D dosen't looks corect
A. Reset the NTP server to resolve any synchronization issues for all devices -- We have an issue only with 1 WLC
B. Check and resolve reachability between Cisco DNA Center and the NTP server We have an issue only with 1 WLC
D. Check and configure NTP on the WLC and synchronize with Cisco DNA Center --- NTP is ""configured"" across the ""network infrastructure"" and Cisco DNA Center.

upvoted 3 times

JOKERR 2 years ago

D is correct.

Excessive time lag between Cisco DNA Center and device: The time difference between Cisco DNA Center and the device IP Address has drifted too far apart. CiscoDNA Center cannot process the device data accurately if the time difference is more than 3 minutes.
Reference: https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center-assurance/1-2- 10/b_cisco_dna_assurance_1_2_10_ug/b_cisco_dna_assurance_1_2_10_ug_chapter_01101.html

upvoted 2 times

[Removed] 1 year, 11 months ago

Well the question stated that NTP is already configured across the network infrastructure as well as DNAC so I think that eliminates D.

upvoted 1 times

Question #188 Topic 1

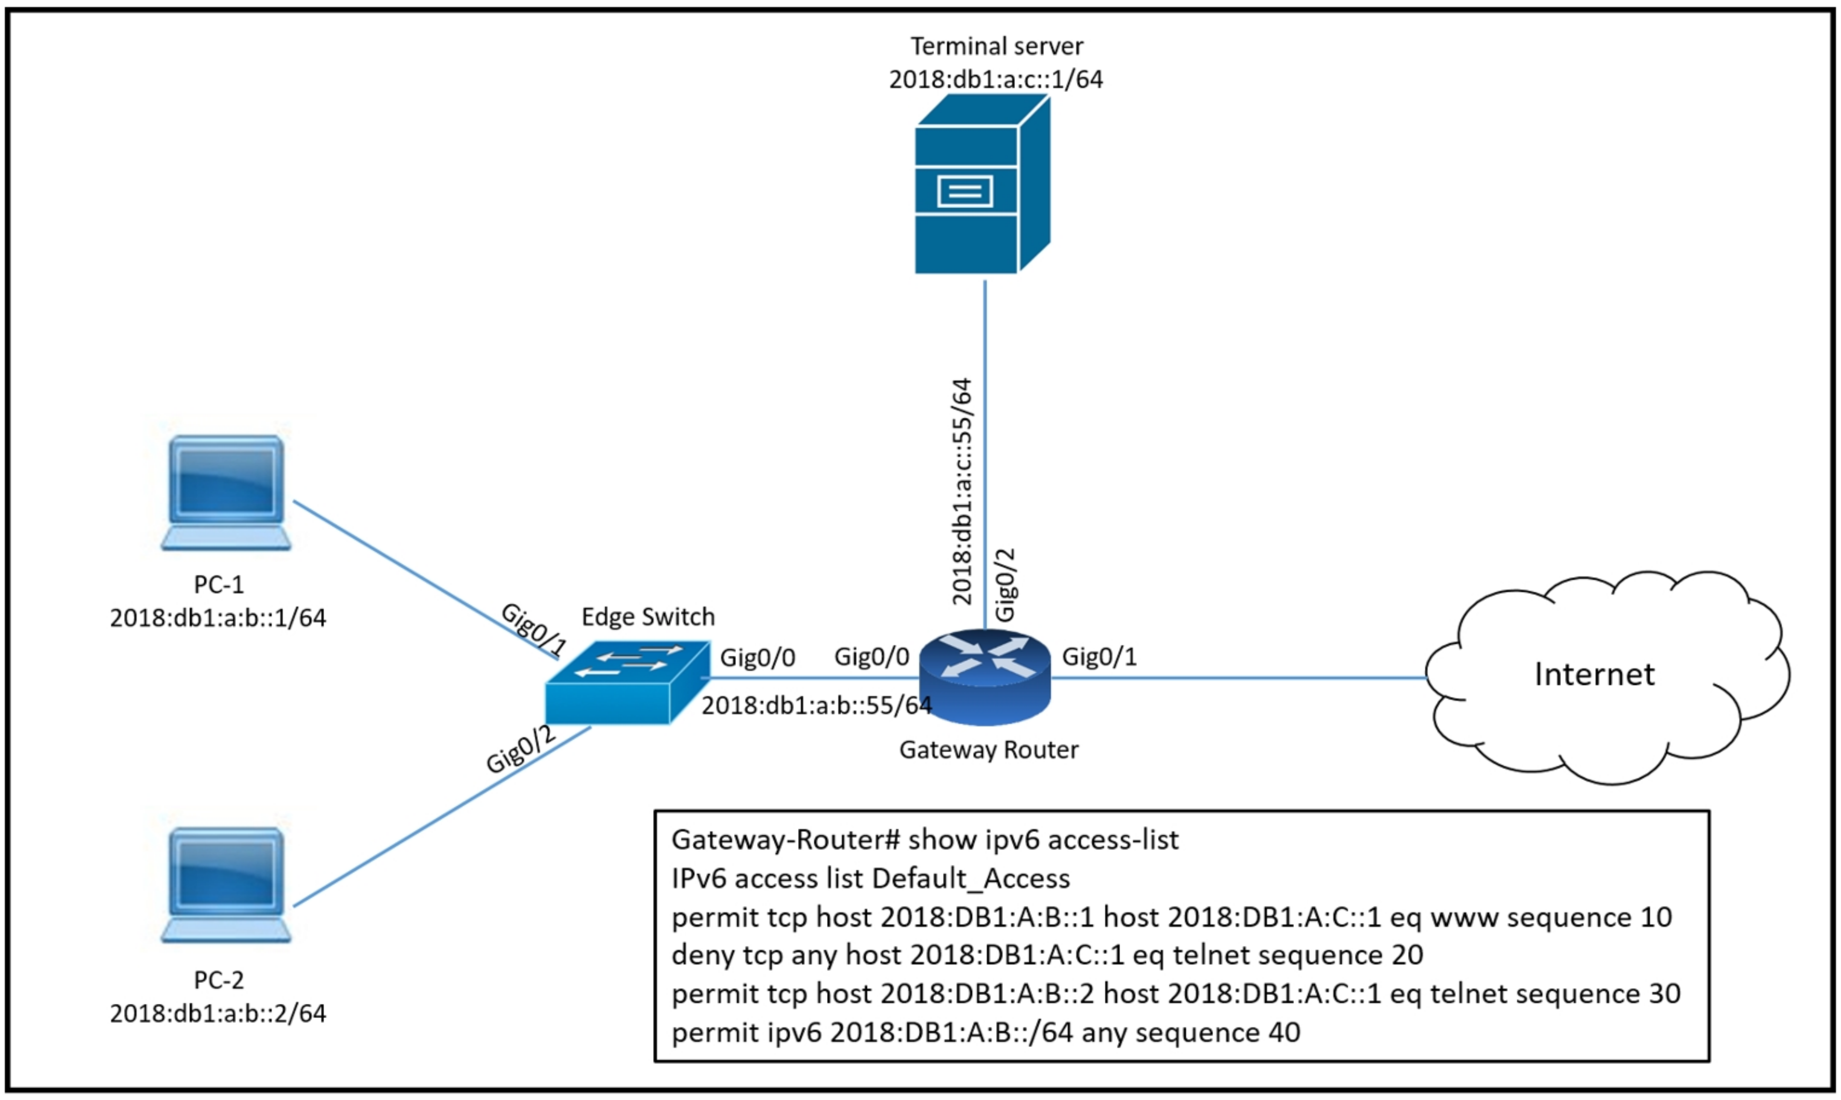

```
Gateway-Router# show ipv6 access-list
IPv6 access list Default_Access
permit tcp host 2018:DB1:A:B::1 host 2018:DB1:A:C::1 eq www sequence 10
deny tcp any host 2018:DB1:A:C::1 eq telnet sequence 20
permit tcp host 2018:DB1:A:B::2 host 2018:DB1:A:C::1 eq telnet sequence 30
permit ipv6 2018:DB1:A:B::/64 any sequence 40
```

Refer to the exhibit. PC-2 failed to establish a Telnet connection to the terminal server.
Which configuration resolves the issue?

A. Gateway-Router(config)#ipv6 access-list Default_Access Gateway-Router(config-ipv6-acl)#sequence 25 permit tcp host 2018:DB1:A:B::2 host 2018:DB1:A:C::1 eq telnet

B. Gateway-Router(config)#ipv6 access-list Default_Access Gateway-Router(config-ipv6-acl)#no sequence 20 Gateway-Router(config-ipv6-acl)#sequence 5 permit tcp host 2018:DB1:A:B::2 host 2018:DB1:A:C::1 eq telnet

C. Gateway-Router(config)#ipv6 access-list Default_Access Gateway-Router(config-ipv6-acl)#permit tcp host 2018:DB1:A:B::2 host 2018:DB1:A:C::1 eq telnet

D. Gateway-Router(config)#ipv6 access-list Default_Access Gateway-Router(config-ipv6-acl)#sequence 15 permit tcp host 2018:DB1:A:B::2 host 2018:DB1:A:C::1 eq telnet

Correct Answer: *D*

Community vote distribution

B (50%) D (50%)

Surfside92 Highly Voted 2 years, 1 month ago

Agree with Amgue that connectivity should already work as pc-2 hits the sequence 30 ACE and as it does not match sequence 10 or 20
There may be a typo in the graphic and sequence 20 should actually read :
deny tcp any host 2018:DB1:A:C::1 eq telnet sequence 20
That would make answer D correct.

However if there's no typo I go for answer B - it tidies things up the most - not completely as sequence 30 remains - but it looks the best fit.

upvoted 5 times

Surfside92 1 year, 11 months ago

Just to update my comment. If you look at the comment below from JOKERR. There is almost certainly a typo in the question above.
That would make the corect answer = D

upvoted 2 times

[Removed] 1 year, 11 months ago

It matches sequence 20 (any) so its getting dropped...

upvoted 4 times

asans Most Recent 3 days, 3 hours ago

Selected Answer: B

Sequence 15 in Answer D "sequence 15 permit tcp host 2018:DB1:A:B::2 host 2018:DB1:A:C::1 eq telnet" is similar to Seq 30 and so the router will just take the accept the ACE but not change the configs on the Default_Access ACL. So D doesnt change anything and thus incorrect

Sequence 5 in Answer B "sequence 5 permit tcp host 2018:DB1:A:B::2 host 2018:DB1:A:C::1 eq telnet" is the same again as Sequence 30 and has the same effect i.e., it doesn't change anything regarding the configs of Default_Access ACL. However the "no sequence 20 " part in Answer B makes the difference. This is what removes the restriction and thus allow Sequence 30 to allow access. Correct answer is B

upvoted 1 times

Chiaretta 5 months ago

The question is where is this ACL applied??? In? Out?

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

optio corret D: because:
rule 5 is duplicated with rule 30. Its not is necesary create rule 5

upvoted 1 times

Dacusai 7 months, 3 weeks ago

Seq 20 block all telnet connection to the server, so we need to introduce one statement before Seq 20 to allow pc2 to access the server.

upvoted 3 times

Nhan 1 year, 6 months ago

the given answer is correct since the acl sequence 10 is permitting the pc-1, then we need to add in a nother permit for pc2 with sequence 15 or 12 or 12 ...

upvoted 1 times

studybuddy10 2 years, 1 month ago

going for D, B violates security. The purpose of this ACL seems to be protection of telnet only as it allows all at seq 40 from those ranges. So only D, they should remove seq 30 though for cleanup.

upvoted 4 times

amgue 2 years, 1 month ago

I think the answer already existe in the show, the permit sequence 30

upvoted 1 times

rob899 3 months, 3 weeks ago

Although the sequence 30 is a good rule to permit PC-2 to Telnet to the server, it is being blocked by the earlier sequence 20 rule which denies ALL telnet traffic to the server.

upvoted 1 times

C_Tw21 2 years, 1 month ago

Hi,
D works ,.
But B should be fine as well.
??

upvoted 1 times

AliMo123 2 years, 1 month ago

it works if we delete sequence 20 but since we have " no sequence 20" in B then only D works here

upvoted 1 times

JOKERR 2 years ago

D is correct.

upvoted 1 times

JOKERR 2 years ago

B also works but is not the best answer. Because removing 20 opens up the ACL to any one to telnet to destination. Adding sequence before it completes our objective while still blocking unwanted devices from accessing via telnet.

upvoted 1 times

JOKERR 2 years ago

Now I am confused, because in exam, I remember that sequence 20 is:
deny tcp any host 2018:db1:A:C::1 eq telnet sequence 20
Which is blocking any telnet connection to terminal server. But here if the seq 20 is not terminal server, either A or D should be correct.

upvoted 1 times

YaPet 1 year, 10 months ago

Probably, here is a mistake for the access-list output

upvoted 1 times

tefacert 1 year, 7 months ago

Hi did you pass the exam, where did you study from ?

upvoted 1 times

Question #189 *Topic 1*

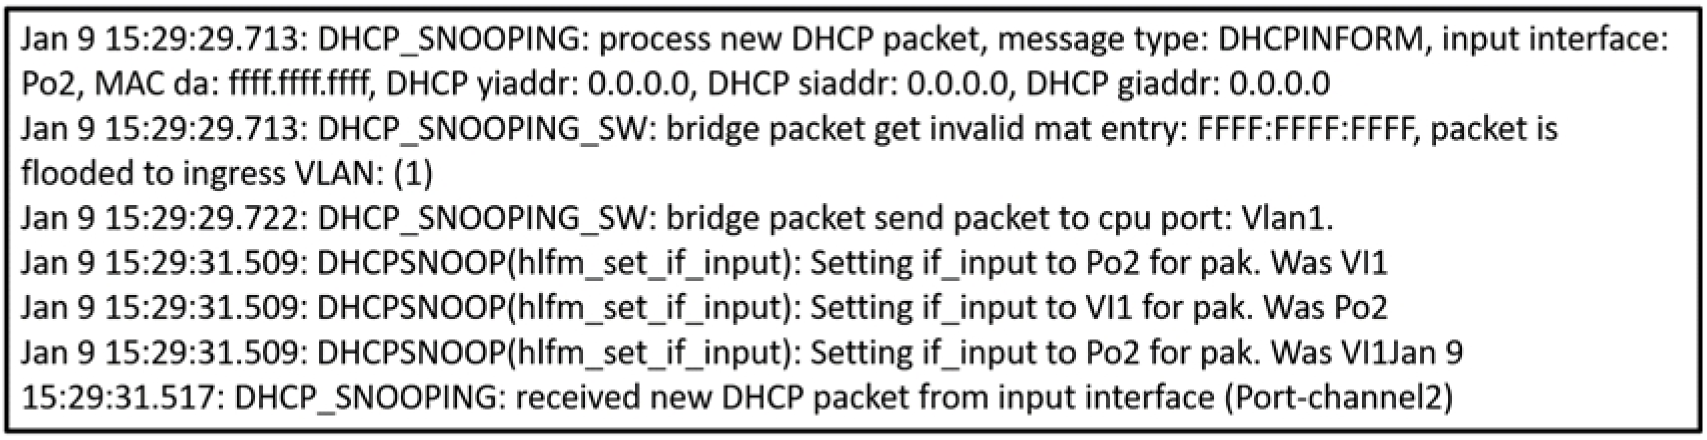

```
Jan 9 15:29:29.713: DHCP_SNOOPING: process new DHCP packet, message type: DHCPINFORM, input interface:
Po2, MAC da: ffff.ffff.ffff, DHCP yiaddr: 0.0.0.0, DHCP siaddr: 0.0.0.0, DHCP giaddr: 0.0.0.0
Jan 9 15:29:29.713: DHCP_SNOOPING_SW: bridge packet get invalid mat entry: FFFF:FFFF:FFFF, packet is
flooded to ingress VLAN: (1)
Jan 9 15:29:29.722: DHCP_SNOOPING_SW: bridge packet send packet to cpu port: Vlan1.
Jan 9 15:29:31.509: DHCPSNOOP(hlfm_set_if_input): Setting if_input to Po2 for pak. Was Vl1
Jan 9 15:29:31.509: DHCPSNOOP(hlfm_set_if_input): Setting if_input to Vl1 for pak. Was Po2
Jan 9 15:29:31.509: DHCPSNOOP(hlfm_set_if_input): Setting if_input to Po2 for pak. Was Vl1Jan 9
15:29:31.517: DHCP_SNOOPING: received new DHCP packet from input interface (Port-channel2)
```

Refer to the exhibit. A network administrator enables DHCP snooping on the Cisco Catalyst 3750-X switch and configures the uplink port (Port-channel2) as a trusted port. Clients are not receiving an IP address, but when DHCP snooping is disabled, clients start receiving IP addresses. Which global command resolves the issue?

A. ip dhcp relay information trust portchannel2

B. ip dhcp snooping

C. ip dhcp snooping trust

D. no ip dhcp snooping information option

Correct Answer: *D*

Reference:

https://community.cisco.com/t5/switching/dhcp-snooping-clients-not-getting-ip-address/td-p/1749969

Community vote distribution

D (100%)

SAMAKEMM 2 months, 2 weeks ago

Selected Answer: D

D is the most relevent

upvoted 1 times

Stylar 6 months, 2 weeks ago

time to re-learn and strengthen dhcp guys right ? :))))

upvoted 2 times

HungarianDish 7 months, 3 weeks ago

Selected Answer: D

This one is the winner:
https://www.kareemccie.com/2016/11/why-do-we-need-ip-dhcp-relay.html

upvoted 3 times

Question #190 *Topic 1*

A customer reports to the support desk that they cannot print from their PC to the local printer id:123456789.
Which tool must be used to diagnose the issue using Cisco DNA Center Assurance?

A. device trace

B. ACL trace

C. path trace

D. application trace

Correct Answer: *C*

Community vote distribution

C (100%)

⊟ **HungarianDish** 7 months, 3 weeks ago

Selected Answer: C

The scenario is from here:
https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2019/pdf/BRKSDN-2426.pdf

upvoted 1 times

⊟ **YaPet** 1 year, 10 months ago

C is correct
https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center-assurance/1-3/b_cisco_dna_assurance_1_3_ug/b_cisco_dna_assurance_1_3_ug_chapter_0111.html

upvoted 1 times

Question #191 *Topic 1*

An engineer configured SNMP notifications sent to the management server using authentication and encrypting data with DES. An error in the response PDU is received as "UNKNOWNUSERNAME, WRONGDIGEST".
Which action resolves the issue?

A. Configure the correct authentication password using SNMPv3 authNoPriv.

B. Configure correct authentication and privacy passwords using SNMPv3 authPriv.

C. Configure correct authentication and privacy passwords using SNMPv3 authNoPriv.

D. Configure the correct authentication password using SNMPv3 authPriv.

Correct Answer: *B*

Community vote distribution

B (100%)

HungarianDish 6 months, 3 weeks ago

Selected Answer: B

Both cases (A, D) generate the same error message.
Still, we need to make sure that both authentication and encryption passwords are correct, otherwise further errors occur.
So for me it's B.

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Data Encryption Standard (DES) is applicable for authPriv only. It narrows down the answer to B or D. B is more precise.

upvoted 1 times

Nonono 1 year, 10 months ago

Selected Answer: B

B correct

upvoted 1 times

leecharxos 1 year, 11 months ago

B is correct, check Table 2 Cisco-Specific Error Messages for SNMPv3 -> authNoPriv
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/configuration/xe-3se/3850/snmp-xe-3se-3850-book/nm-snmp-snmpv3.html

upvoted 2 times

wts 1 year, 9 months ago

From the table, you can find out that these messages correspond to
Configured Security Level:
- authNoPriv
- authPriv
and to Security Level of Incoming SNMP Message:
- authNoPriv with incorrect authentication password
- authPriv with incorrect authentication password and correct privacy password.

But how to correct the error and answer correctly is not clear.

upvoted 1 times

wts 1 year, 9 months ago

- DES encryption MEAN THAT the security level of SNMPv3 is authPriv.(tab1)
- Such messages for this level MEAN THAT "authPriv with incorrect authentication password and correct privacy password" OR "authPriv with incorrect authentication password and incorrect privacy password"(tab2)
- So this NEEDS TO "Configure correct authentication and privacy passwords using SNMPv3 authPriv" OR "Configure the correct authentication password using SNMPv3 authPriv". (question)

So... B or D?

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

Question #192 *Topic 1*

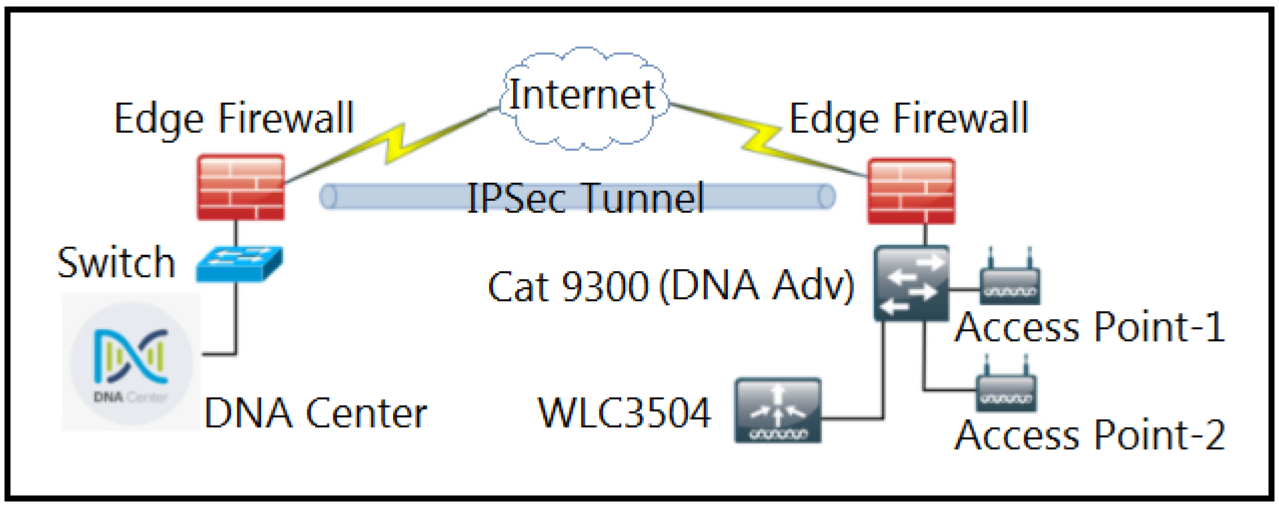

Refer to the exhibit. A network administrator is discovering a Cisco Catalyst 9300 and a Cisco WLC 3504 in Cisco DNA Center. The Catalyst 9300 is added successfully. However, the WLC is showing the error "uncontactable" when the administrator tries to add it in Cisco DNA Center. Which action discovers WLC in Cisco DNA Center successfully?

A. Delete the WLC 3504 from Cisco DNA Center and add it to Cisco DNA Center again.

B. Add the WLC 3504 under the hierarchy of the Catalyst 9300 connected devices.

C. Copy the .cert file from the Cisco DNA Center on the USB and upload it to the WLC 3504.

D. Copy the .pem file from the Cisco DNA Center on the USB and upload it to the WLC 3504.

Correct Answer: *D*

Community vote distribution

D (100%)

[Removed] 4 months ago

From the book
"Cisco DNA Center is a massive topic that is beyond the scope of the ENARSI exam. The official exam objectives for ENARSI state that you should be able to “troubleshoot network problems using Cisco DNA Center Assurance (connectivity, monitoring, device health, net□work health).” Therefore, this section remains focused on this objective."

Cisco... what the actual fuck...

upvoted 3 times

Brand 3 months, 3 weeks ago

In the "Exam Topics" of ENARSI it also says and I quote "The above topics are likely to be included on the 300-410 ENARSI exam. The topics are subject to change at any time bla bla bla"

upvoted 1 times

chris110 3 months, 3 weeks ago

welcome to a never resting industrie ;)

upvoted 3 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

option D is corerct
https://community.cisco.com/t5/cisco-digital-network-architecture-dna/dnac-assurance-wlc3504/td-p/3841805

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

https://community.cisco.com/t5/cisco-digital-network/dnac-assurance-wlc3504/td-p/3841805

upvoted 2 times

Question #193 *Topic 1*

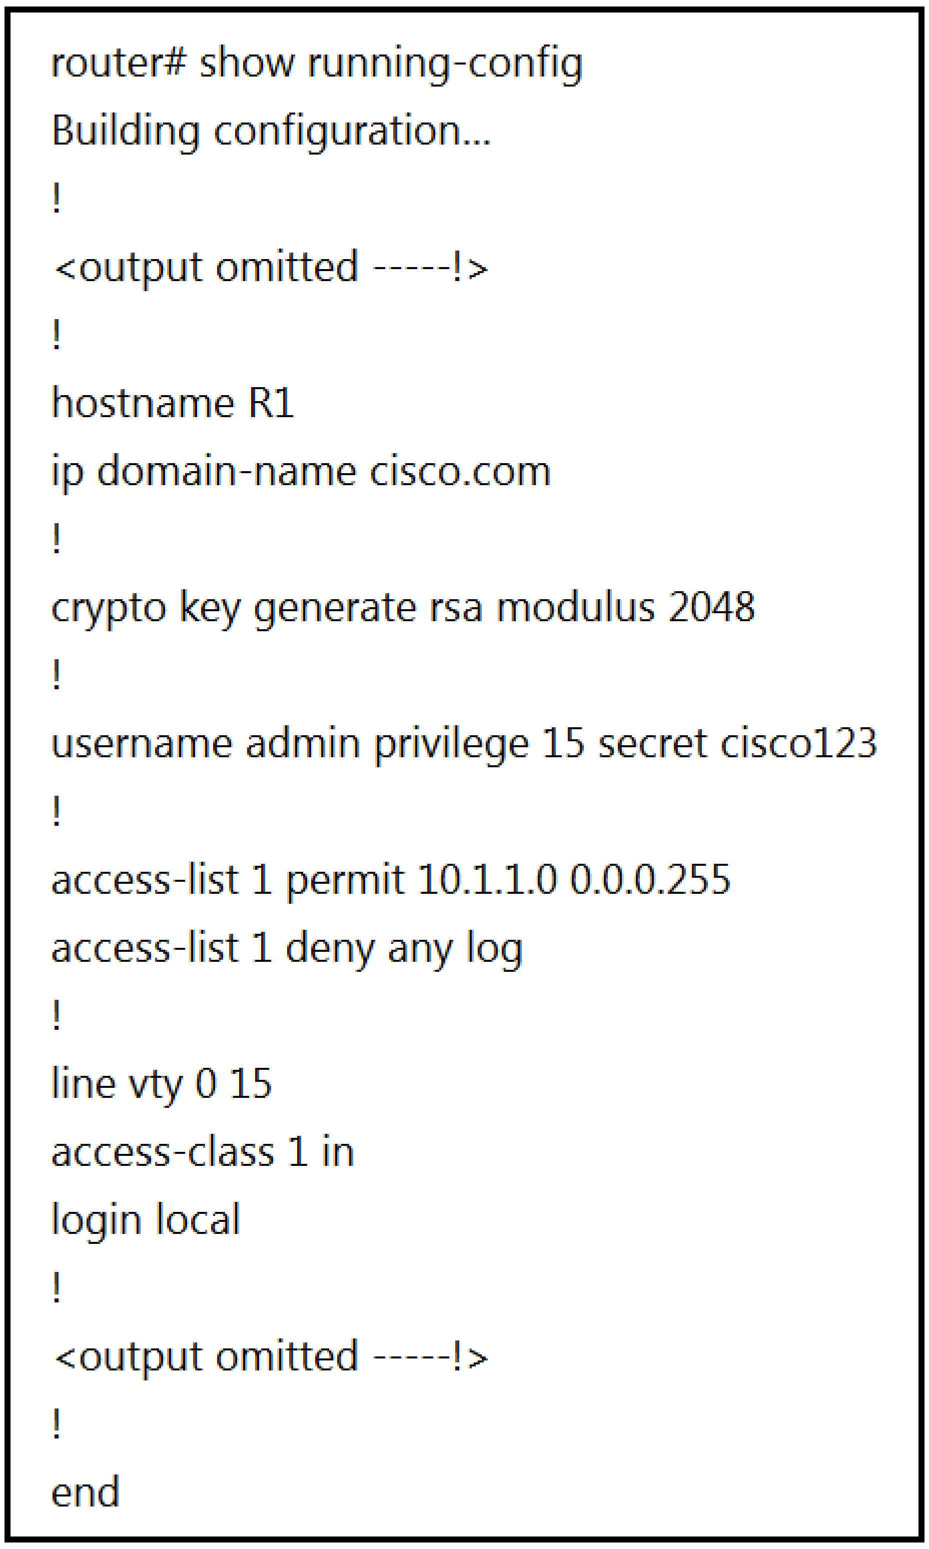

```
router# show running-config
Building configuration...
!
<output omitted -----!>
!
hostname R1
ip domain-name cisco.com
!
crypto key generate rsa modulus 2048
!
username admin privilege 15 secret cisco123
!
access-list 1 permit 10.1.1.0 0.0.0.255
access-list 1 deny any log
!
line vty 0 15
access-class 1 in
login local
!
<output omitted -----!>
!
end
```

Refer to the exhibit. A user cannot SSH to the router.

What action must be taken to resolve this issue?

A. Configure transport input ssh

B. Configure transport output ssh

C. Configure ip ssh version 2

D. Configure ip ssh source-interface loopback0

Correct Answer: *A*

Community vote distribution

A (80%) C (20%)

MasterMatt 8 months, 2 weeks ago

Selected Answer: A

ssh is enabled by default but temporarily disabled if the rsa key is not generated. Once the key is generated, and you have local account plus the transport input ssh you should be able to login with SSH.

upvoted 2 times

ERICKPORRAS 1 year, 3 months ago

Selected Answer: A

A is correct, C is incorrect because:
If you do not enter this command "ip ssh version 1/ 2 " or do not specify a keyword, the SSH server selects the latest SSH version supported by the SSH client. For example, if the SSH client supports SSHv1 and SSHv2, the SSH server selects SSHv2.

check it: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960x/software/15-0_2_EX/security/configuration_guide/b_sec_152ex_2960-x_cg/b_sec_152ex_2960-x_cg_chapter_01001.html

upvoted 2 times

Kimaf 1 year, 8 months ago

Selected Answer: C

Is the correct version of SSH specified? By default, both version 1 and 2 are enabled. However, with the ip ssh version {1 | 2} command, you can change the version to just 1 or 2. If clients are connecting with version 2 and the device is configured for version 1, the SSH connection will fail; the same is true if clients are using version 1 and the devices are configured for version 2. To check the version of SSH that is running, use the show ip ssh command, as shown in Example 23-5. If it states version 1.99, it means versions 1 and 2 are running. If it states version 1, then SSHv1 is running, and if it states version 2, then SSHv2 is running.
Has the correct key size been specified? SSHv2 uses an RSA key size of 768 or greater. If you were using a smaller key size with SSHv1 and then switched to SSHv2, you would need to create a new key with the correct size; otherwise, SSHv2 would not work. If you are using SSHv2 but accidentally specify a key size less than 768, SSHv2 connections are not allowed.
I have based my answer on OCG ENARSI BOOK PAGE 874 and since its specifies 🔑 size of greater than 768.

upvoted 1 times

Surfside92 2 years, 1 month ago

I think the answer = C
The default transport input is both telnet and ssh so that rules out answer A.
The config "ip ssh version 2" is part of the required ssh configuration - and that is missing from the output.

upvoted 1 times

tsabee 2 years, 1 month ago

You've partially right, but the default function was changed:
according to command reference:
"Defaults
No protocols are allowed on the auxiliary (AUX), console, tty, and vty lines.
...
Cisco devices do not accept incoming network connections to tty lines by default. You must specify an incoming transport protocol or specify the transport input all command before the line will accept incoming connections.
...
This behavior is new as of Cisco IOS Release 15.4(3)M4. Previous to Cisco IOS Release 15.4(3)M4, the default was the transport input all command. If you are upgrading to a release later than Cisco IOS Release 15.4(3)M4, you must configure the transport input none command, or you will be locked out of your device."

https://www.cisco.com/c/en/us/td/docs/ios/termserv/command/reference/tsv_book/tsv_s1.html

upvoted 5 times

myrmike 1 year, 11 months ago

To add on if a crypto key is generated the ssh version 1.99 is enabled.

upvoted 3 times

tsabee 2 years, 1 month ago

So I think the correct answer is A.

upvoted 3 times

[Removed] 1 year, 11 months ago

Nope... If a ssh version isnt specified, the latest version of ssh is selected.

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960x/software/15-0_2_EX/security/configuration_guide/b_sec_152ex_2960-x_cg/b_sec_152ex_2960-x_cg_chapter_01001.html

upvoted 4 times

OakA1 2 years, 2 months ago

I don't see any of the answers being correct. The default transport input is both telnet and ssh. Everything is enabled for SSH: domain and crypto key... There is also a local user configured. For me the only way a user can't login if he or she is connecting from a subnet that is not specified in the ACL.

upvoted 1 times

JOKERR 2 years ago

Default transport is none. You have to specify explicitly which protocol you want to allow. Otherwise you will get this:

```
ER1#telnet 172.16.45.1
Trying 172.16.45.1 ...
% Connection refused by remote host
ER1#
ER1#ssh -l admin 172.16.45.1
% Connection refused by remote host
```

upvoted 4 times

examShark 2 years, 4 months ago

Te given answer is correct

upvoted 2 times

Abudi 1 year, 1 month ago

there is an evidence here that you are actually typing “The given answer is correct” in each question and not copy/pasting it xD

upvoted 3 times

Question #194 Topic 1

An engineer configured a Cisco router to send reliable and encrypted notifications for any events to the management server. It was noticed that the notification messages are reliable but not encrypted.
Which action resolves the issue?

A. Configure all devices for SNMPv3 informs with auth.

B. Configure all devices for SNMPv3 informs with priv.

C. Configure all devices for SNMPv3 traps with auth.

D. Configure all devices for SNMPv3 traps with priv.

Correct Answer: *B*

Community vote distribution

B (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

Option B:

https://community.microfocus.com/it_ops_mgt/nom/f/nom-user-discussions/81954/nnmi-support-tip-difference-between-snmp-inform-and-snmp-trap-from-nnmi-point-of-view

upvoted 1 times

Titini 10 months, 1 week ago

B The "inform" message is a type of SNMPv3 notification that requires an acknowledgment from the recipient device, making it a reliable way of transmitting information. The "priv" option in SNMPv3 provides encryption of the data being transmitted, ensuring that the information is secure and cannot be intercepted by unauthorized users.

upvoted 3 times

wts 1 year, 9 months ago

Selected Answer: B

B is correct

upvoted 1 times

Carl1999 1 year, 10 months ago

Selected Answer: B

B is correct.
"informs" is "reliable"

upvoted 3 times

toni2 1 year, 10 months ago

I think B i correct:
The following example shows how to configure a remote user to receive traps at the “noAuthNoPriv” security level when the SNMPv3 security model is enabled:

```
Device(config)# snmp-server group group1 v3 noauth
Device(config)# snmp-server user remoteuser1 group1 remote 10.12.8.4
Device(config)# snmp-server host 10.12.8.4 informs version 3 noauth remoteuser config
```

upvoted 1 times

steiger 2 years ago

Selected Answer: B

B is my choice

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

error_909 2 years, 3 months ago

The major difference between an inform request and a trap is that an SNMP agent has no way of knowing if an SNMP trap was received by the SNMP manager. However, an SNMP inform request packet will be sent continually until the sending SNMP manager receives an SNMP acknowledgement.

"Reliable" is stated clearly in the question

upvoted 6 times

jarz 1 year, 1 month ago

Nice, thanks for pointing this out!

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

RHK0783 2 years, 7 months ago

The router configured to "send" out something is "TRAP". Auth makes it reliable and Priv makes it confidential ...
Answer is D.

upvoted 1 times

frzzt123 2 years, 7 months ago

Not true, reliable means he waits for a response from the NMS. Traps are not reliable. Informs are.

upvoted 3 times

CiscoCCNPDream 2 years, 7 months ago

The answer should be B because the major difference between an inform request and a trap is that an SNMP agent has no way of knowing if an SNMP trap was received by the SNMP manager. However, an SNMP inform request packet will be sent continually until the sending SNMP manager receives an SNMP acknowledgement. The keyword is "reliable" here

upvoted 1 times

oasc 2 years, 8 months ago

Actually B is right since Inform packets are performed for events in SNMP, while traps are independent sents from agent to manager. But I guess the trick is in the word event

upvoted 1 times

oasc 2 years, 8 months ago

sorry D

upvoted 1 times

oasc 2 years, 8 months ago

should not be B

upvoted 1 times

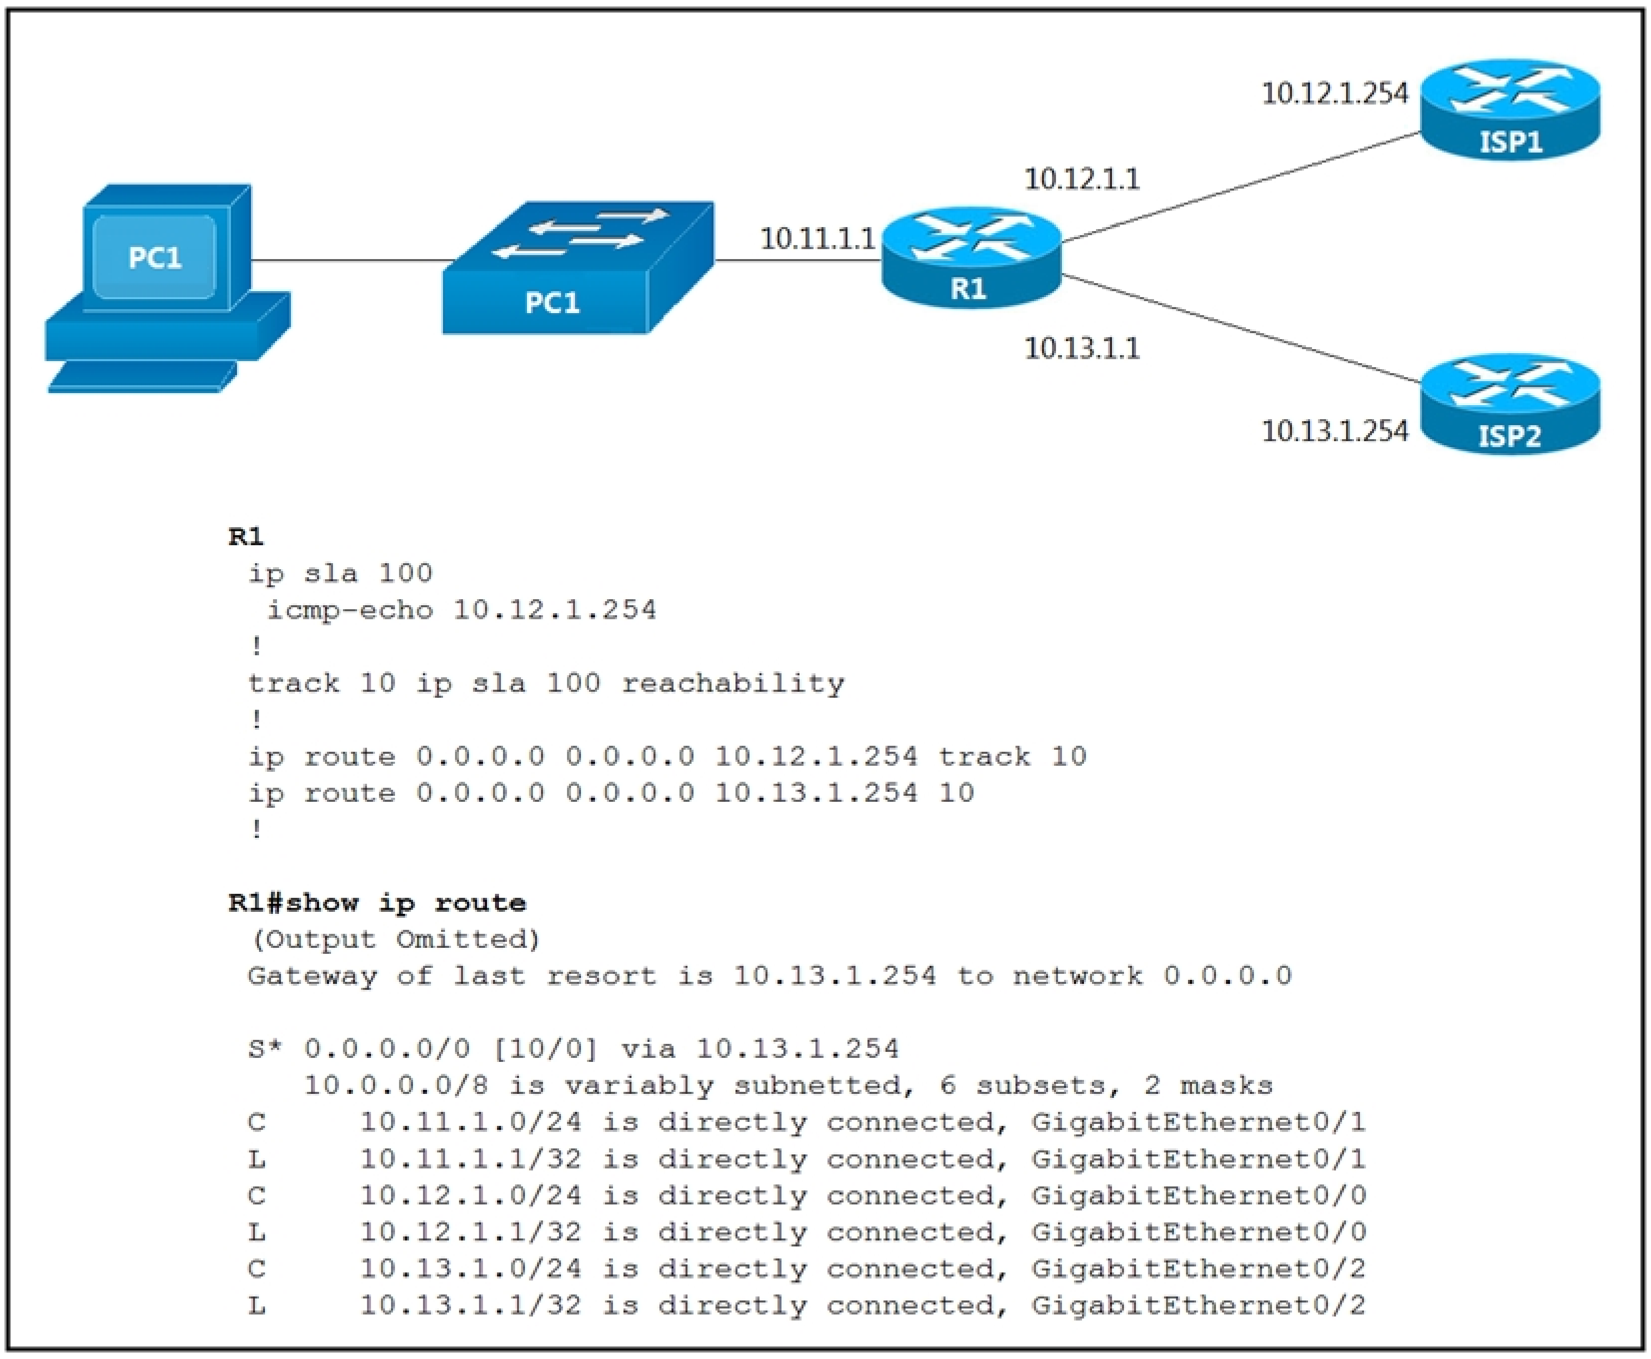

```
R1
 ip sla 100
  icmp-echo 10.12.1.254
 !
 track 10 ip sla 100 reachability
 !
 ip route 0.0.0.0 0.0.0.0 10.12.1.254 track 10
 ip route 0.0.0.0 0.0.0.0 10.13.1.254 10
 !

R1#show ip route
 (Output Omitted)
 Gateway of last resort is 10.13.1.254 to network 0.0.0.0

 S* 0.0.0.0/0 [10/0] via 10.13.1.254
    10.0.0.0/8 is variably subnetted, 6 subsets, 2 masks
 C     10.11.1.0/24 is directly connected, GigabitEthernet0/1
 L     10.11.1.1/32 is directly connected, GigabitEthernet0/1
 C     10.12.1.0/24 is directly connected, GigabitEthernet0/0
 L     10.12.1.1/32 is directly connected, GigabitEthernet0/0
 C     10.13.1.0/24 is directly connected, GigabitEthernet0/2
 L     10.13.1.1/32 is directly connected, GigabitEthernet0/2
```

Refer to the exhibit. An engineer is monitoring reachability of the configured default routes to ISP1 and ISP2. The default route from ISP1 is preferred if available.

How is this issue resolved?

A. Use the icmp-echo command to track both default routes.

B. Use the same AD for both default routes.

C. Start IP SLA by matching numbers for track and ip sla commands.

D. Start IP SLA by defining frequency and scheduling it.

Correct Answer: *D*

Community vote distribution

D (100%)

NoUserName1234 1 year, 3 months ago

Bad question ISP 1 is preferred, but routes go to ISP2 .
Nevertheless, provided answer in context is probably right

upvoted 3 times

Emery12 11 months, 2 weeks ago

That's the whole point, it's not working as expected! For that to happen, the ip sla should start first using scheduling

upvoted 1 times

WAKIDI 1 year, 5 months ago

Selected Answer: D

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/configuration/15-mt/sla-15-mt-book/sla_icmp_echo.html

upvoted 1 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

Question #196 Topic 1

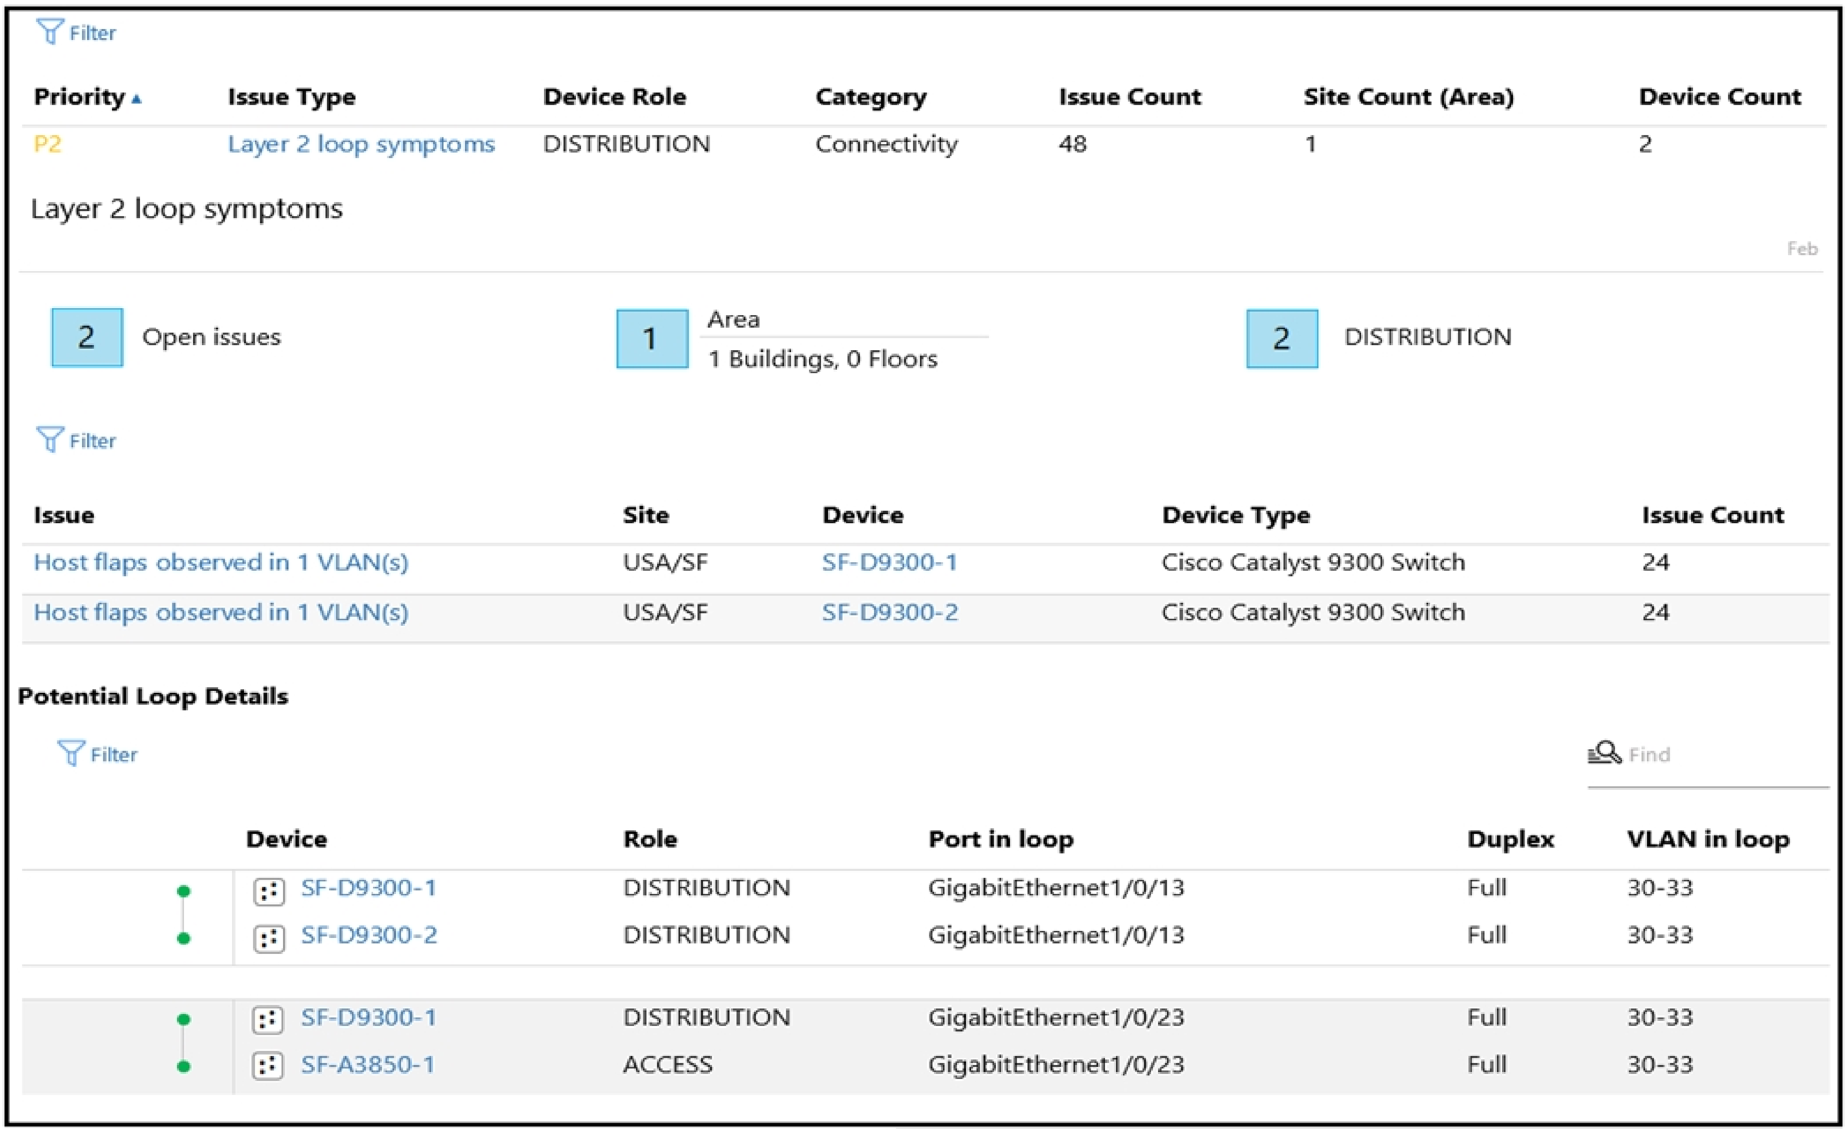

Filter

Priority ▲	Issue Type	Device Role	Category	Issue Count	Site Count (Area)	Device Count
P2	Layer 2 loop symptoms	DISTRIBUTION	Connectivity	48	1	2

Layer 2 loop symptoms

Feb

2 Open issues

1 Area
1 Buildings, 0 Floors

2 DISTRIBUTION

Filter

Issue	Site	Device	Device Type	Issue Count
Host flaps observed in 1 VLAN(s)	USA/SF	SF-D9300-1	Cisco Catalyst 9300 Switch	24
Host flaps observed in 1 VLAN(s)	USA/SF	SF-D9300-2	Cisco Catalyst 9300 Switch	24

Potential Loop Details

Filter Find

Device	Role	Port in loop	Duplex	VLAN in loop
SF-D9300-1	DISTRIBUTION	GigabitEthernet1/0/13	Full	30-33
SF-D9300-2	DISTRIBUTION	GigabitEthernet1/0/13	Full	30-33
SF-D9300-1	DISTRIBUTION	GigabitEthernet1/0/23	Full	30-33
SF-A3850-1	ACCESS	GigabitEthernet1/0/23	Full	30-33

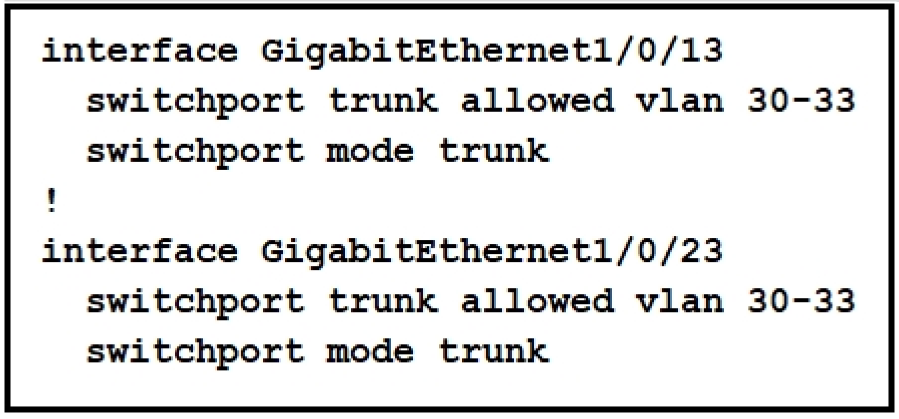

```
interface GigabitEthernet1/0/13
  switchport trunk allowed vlan 30-33
  switchport mode trunk
!
interface GigabitEthernet1/0/23
  switchport trunk allowed vlan 30-33
  switchport mode trunk
```

Refer to the exhibits. An engineer identified a Layer 2 loop using DNAC. Which command fixes the problem in the SF-D9300-1 switch?

A. spanning-tree portfast bpduguard

B. no spanning-tree uplinkfast

C. spanning-tree backbonefast

D. spanning-tree loopguard default

Correct Answer: *A*

Community vote distribution

D (80%) A (20%)

OakA1 Highly Voted 2 years, 2 months ago

The answer should be D. The A enables bpduguard on access ports. We have trunks here. So, loopguard enabled on the trunks will solve the issue. https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/10596-84.html give a good explanation

upvoted 12 times

AliMo123 Highly Voted 2 years, 2 months ago

wrong answer
portfast is only configured on access port not trunk, so D is correct

upvoted 6 times

[Removed] 1 year, 11 months ago

Im assuming the answer is A because there is a access port connected to a trunk port on that gig interface which would cause loops to occur

upvoted 1 times

Carl1999 1 year, 10 months ago

Where do you know the access port?
"Role access" means access switch not access port.
upvoted 2 times

BTK0311 Most Recent 3 months ago

Selected Answer: A

Enabling "spanning-tree portfast bpduguard" on access ports can help prevent Layer 2 loops by shutting down the port if a BPDU (Bridge Protocol Data Unit) is received on the port. This is a common best practice to ensure that access ports do not participate in creating loops.Enabling "spanning-tree loopguard default" globally on a switch will activate the loop guard feature on all designated ports. Loop guard is used to prevent Layer 2 loops in spanning tree networks by detecting and responding to BPDUs (Bridge Protocol Data Units) that are not received as expected.

However, enabling "spanning-tree loopguard default" across all designated ports may not be the most appropriate action in all situations. It's a broad change that can affect the entire switch, potentially leading to unwanted consequences in certain network setups.
upvoted 1 times

ridonak230 3 months ago

Selected Answer: D

Answer D is the correct one !

D. spanning-tree loopguard default
upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: D

As I understand, there are some cases, when we would enable portfast and bpduguard on trunk links (for instance, when connecting to ESXi server). Good thread:
https://community.cisco.com/t5/switching/enable-bpduguard-on-spanning-tree-portfast-trunk-port-yes-or-no/td-p/2534826

Based on the output, these are two switches that are connected through the affected trunk ports. So, I find loopguard to be the appropriate solution.
https://networklessons.com/spanning-tree/spanning-tree-loopguard-udld
upvoted 2 times

Brand 3 months, 3 weeks ago

We do that but the portfast command requires "trunk" option. It's not the case for A.
upvoted 1 times

Slinky 8 months, 2 weeks ago

Selected Answer: D

Configuring BPDUguard here is just going to shut down both trunks and there will be no traffic.
upvoted 3 times

BECAUSE 1 year, 2 months ago

Selected answer is correct.
- Not configured under the interface.
- https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/10586-65.html
upvoted 1 times

cyrus777 1 year, 8 months ago

A
https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/10596-84.html
upvoted 3 times

cyrus777 1 year, 8 months ago

The loop guard feature is enabled on a per-port basis. However, as long as it blocks the port on the STP level, loop guard blocks inconsistent ports on a per-VLAN basis (because of per-VLAN STP). That is, if BPDUs are not received on the trunk port for only one particular VLAN, only that VLAN is blocked (moved to loop-inconsistent STP state). For the same reason, if enabled on an EtherChannel interface, the entire channel is blocked for a particular VLAN, not just one link (because EtherChannel is regarded as one logical port from the STP point of view).

On which ports should the loop guard be enabled? The most obvious answer is on the blocking ports. However, this is not totally correct. Loop guard must be enabled on the non-designated ports (more precisely, on root and alternate ports) for all possible combinations of active topologies. As long as the loop guard is not a per-VLAN feature, the same (trunk) port might be designated for one VLAN and non-designated for the other. The possible failover scenarios should also be taken into account.
upvoted 1 times

cyrus777 1 year, 8 months ago

Understanding BPDU Guard
The BPDU guard feature can be globally enabled on the switch or can be enabled per interface, but the feature operates with some differences.
upvoted 1 times

cyrus777 1 year, 8 months ago

At the global level, you enable BPDU guard on Port Fast-enabled STP ports by using the spanning-tree portfast bpduguard default global configuration command. Spanning tree shuts down STP ports that are in a Port Fast-operational state if any BPDU is received on those ports. In a valid configuration, Port Fast-enabled STP ports do not receive BPDUs. Receiving a BPDU on a Port Fast-enabled port signals an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature puts the interface in the error-disabled state.

upvoted 1 times

cyrus777 1 year, 8 months ago

At the interface level, you enable BPDU guard on any STP port by using the spanning-tree bpduguard enable interface configuration command without also enabling the Port Fast feature. When the STP port receives a BPDU, it is put in the error-disabled state.
The BPDU guard feature provides a secure response to invalid configurations because you must manually put the interface back in service. Use the BPDU guard feature in a service-provider network to prevent an access port from participating in the spanning tree.
You can enable the BPDU guard feature for the entire switch or for an interface.

upvoted 1 times

bogd 1 year, 10 months ago

Selected Answer: D

Loopguard

upvoted 2 times

Carl1999 1 year, 10 months ago

Selected Answer: A

```
Switch(config)#int gi0/1
Switch(config-if)#switchport mode trunk
Switch(config-if)#spanning-tree portfast
%Warning: portfast should only be enabled on ports connected to a single
host. Connecting hubs, concentrators, switches, bridges, etc... to this
interface when portfast is enabled, can cause temporary bridging loops.
Use with CAUTION
%Portfast has been configured on GigabitEthernet0/1 but will only
have effect when the interface is in a non-trunking mode.
Switch(config-if)#
Switch(config-if)# spanning-tree portfast trunk
%Warning: portfast should only be enabled on ports connected to a single
host. Connecting hubs, concentrators, switches, bridges, etc... to this
interface when portfast is enabled, can cause temporary bridging loops.
Use with CAUTION

Switch(config-if)#exit
Switch(config)#spanning-tree portfast bpduguard default
```

upvoted 1 times

Question #197 *Topic 1*

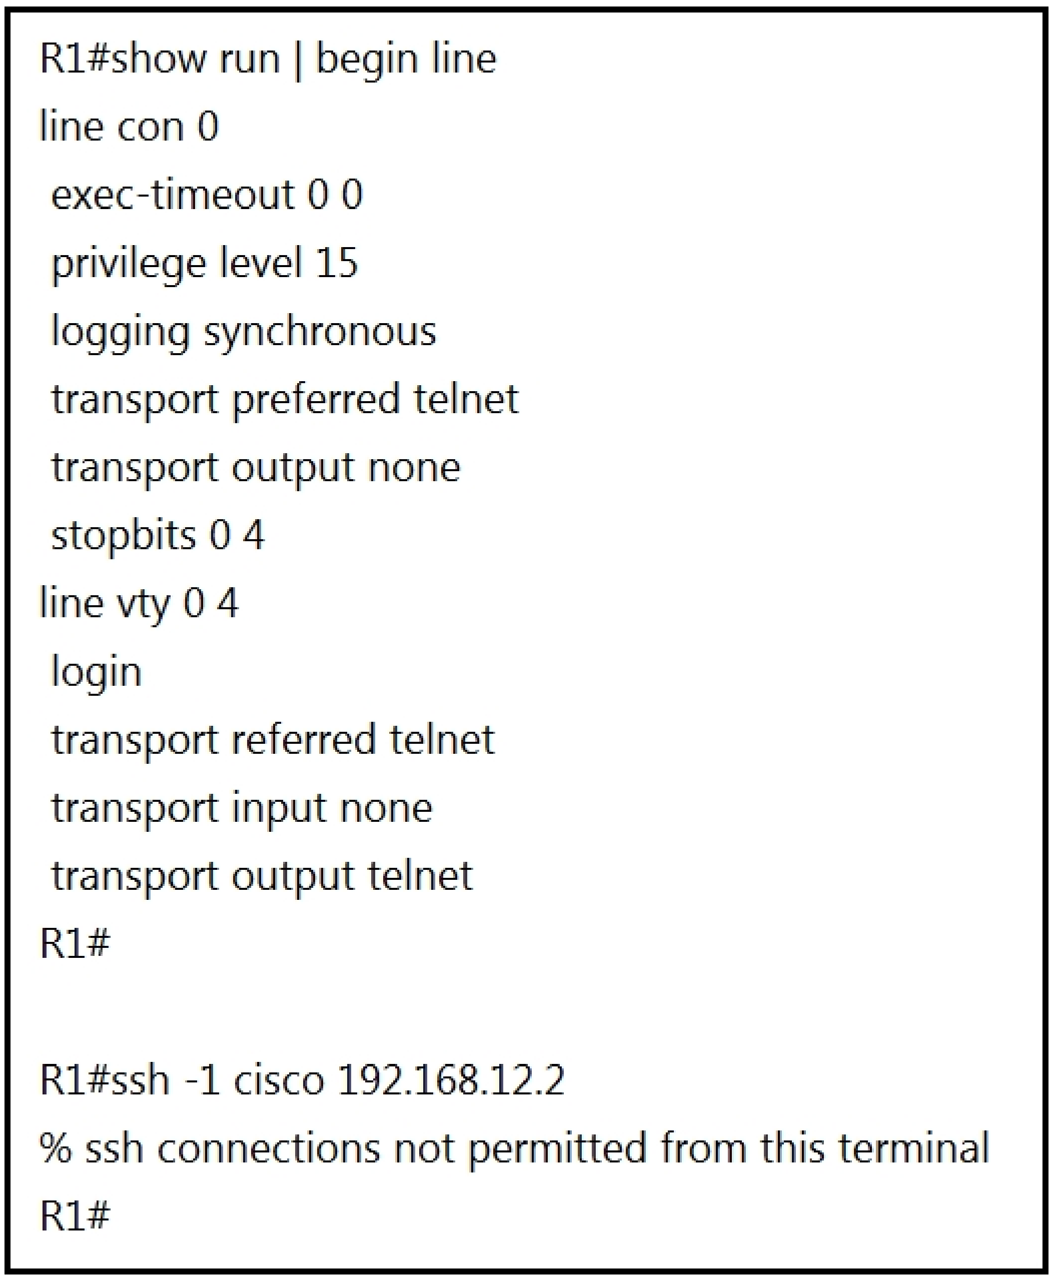

```
R1#show run | begin line
line con 0
 exec-timeout 0 0
 privilege level 15
 logging synchronous
 transport preferred telnet
 transport output none
 stopbits 0 4
line vty 0 4
 login
 transport referred telnet
 transport input none
 transport output telnet
R1#

R1#ssh -1 cisco 192.168.12.2
% ssh connections not permitted from this terminal
R1#
```

Refer to the exhibit. An engineer receives this error message when trying to access another router in-band from the serial interface connected to the console of
R1.
Which configuration is needed on R1 to resolve this issue?

A. R1(config)#line vty 0 R1(config-line)# transport output ssh

B. R1(config)#line console 0 R1(config-line)# transport output ssh

C. R1(config)#line console 0 R1(config-line)# transport preferred ssh

D. R1(config)#line vty 0 R1(config-line)# transport output ssh R1(config-line)# transport preferred ssh

Correct Answer: *D*

Community vote distribution

B (89%) 11%

tseen Highly Voted 10 months, 2 weeks ago

Selected Answer: B

```
R2#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)#lin vty 0
R2(config-line)#transport output ssh
R2(config-line)#transport preferred ssh
R2(config-line)#^Z
R2#conf t
*Feb 8 20:47:02.183: %SYS-5-CONFIG_I: Configured from console by console
R2#ssh -l admin 10.0.0.1
% ssh connections not permitted from this terminal
R2#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)#line con 0
R2(config-line)#transport output ssh
R2(config-line)#transport output ssh
R2(config-line)#^Z
R2#ssh -l admin 10.0.0.1
```

```
*Feb 8 20:47:37.523: %SYS-5-CONFIG_I: Configured from console by console
R2#ssh -l admin 10.0.0.1
Password:
R1#
```

upvoted 5 times

inteldarvid Most Recent 5 months, 2 weeks ago

B option corerct

upvoted 1 times

guy276465281819372 5 months, 3 weeks ago

Selected Answer: B

b is correct

upvoted 1 times

HungarianDish 7 months, 3 weeks ago

Selected Answer: B

I also tested it, and the result was answer B for me, too.
On R1 under "line con 0" "transport output none".
When first connecting through serial cable to R1 and then connecting via ssh to the other router "% ssh connections not permitted from this terminal" appeared.
Solution: On R1 under "line con 0" "transport output ssh".
Changing the configuration under "line vty 0 4" had no effect.

upvoted 2 times

HungarianDish 7 months, 3 weeks ago

Also tested the other scenario: first ssh to R1 and then ssh to the other device.
In this case, "line vty 0 4" "transport output none" produced the same error.
"line vty 0 4" "transport output ssh" was required for successful connection to the other device via ssh.

upvoted 1 times

Huntkey 1 year, 2 months ago

Selected Answer: D

I never heard that you could use SSH protocol on the Serial interface... My understanding of the question is that you use serial console to connect to R1 then use SSH to connect to another device over the VTY

Look at this post and look at Aaron's response
https://community.cisco.com/t5/switching/transport-preferred-ssh-command-at-console-line/m-p/4469002#M511050

upvoted 2 times

Huntkey 1 year, 2 months ago

Damn I meant for B...

upvoted 1 times

networkWiz 1 year, 4 months ago

Selected Answer: B

B is the correct answer. it states in the question access from the Serial interface (console cable).

upvoted 2 times

Nhan 1 year, 6 months ago

D is correct answer,

upvoted 1 times

YaPet 1 year, 10 months ago

Selected Answer: B

B is correct

upvoted 3 times

krn007 1 year, 11 months ago

Selected Answer: B

Correct Answer :B

upvoted 4 times

Mr_RaCailum 2 years, 4 months ago

This is the most basic question... B is the answer obviously.

upvoted 3 times

examShark 2 years, 4 months ago

B is the correct answer

upvoted 2 times

RHK0783 2 years, 7 months ago

the error is about terminal connection. not console ...

upvoted 1 times

ZachTL11 2 years, 8 months ago

R1(config)#line console 0
R1(config-line)# transport output ssh

upvoted 3 times

DaanB 2 years, 8 months ago

This article (https://packetu.com/2016/07/07/understanding-transport-output-access-class/) supports my and other people opinion that the answer is C, that you should change the config for the line console 0, not for the line vty 0 4

upvoted 1 times

DaanB 2 years, 8 months ago

Not C, it's B

upvoted 5 times

steiger 2 years, 1 month ago

How do you use ssh when connecting to the console?

upvoted 1 times

Alnet 2 years, 1 month ago

It's in the wording of the question. It says that you (engineer) are connected to this router via it's console. So your current session is under the consoles rules/config. Altering the VTY on the router which your logged into the console will have no effect on you being able to SSH outbound to another device.

upvoted 6 times

Question #198 Topic 1

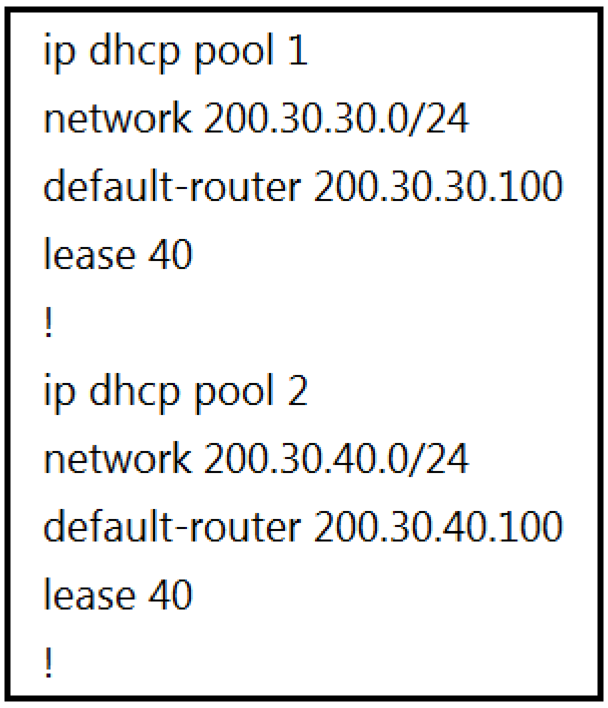

```
ip dhcp pool 1
network 200.30.30.0/24
default-router 200.30.30.100
lease 40
!
ip dhcp pool 2
network 200.30.40.0/24
default-router 200.30.40.100
lease 40
!
```

Refer to the exhibit. The server for the finance department is not reachable consistently on the 200.30.40.0/24 network and after every second month it gets a new
IP address.
What two actions must be taken to resolve this issue? (Choose two.)

A. Configure the server to use DHCP on the network with default gateway 200.30.40.100.

B. Configure the server with a static IP address and default gateway.

C. Configure the router to exclude a server IP address.

D. Configure the server to use DHCP on the network with default gateway 200.30.30.100.

E. Configure the router to exclude a server IP address and default gateway.

Correct Answer: *BC*

Community vote distribution

BE (50%) BC (44%) 6%

Dirkd0344 Highly Voted 2 years ago

The given answer is correct. The default gateway's address is automatically reserved when the DHCP pool is created.
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_dhcp/configuration/12-4/dhcp-12-4-book/config-dhcp-server.html

upvoted 8 times

ookr 1 year, 8 months ago

The issue is that "The IP address configured on the router interface is automatically excluded from the DHCP address pool." but 200.30.40.100 could be on a different router. Also, what if the def.gw. is a HSRP address? The DHCP server and the def. gw. don't have to be the same device. So it could be either BC or BE. I'd say we have to toss a coin.

upvoted 3 times

night_wolf_in Most Recent 1 month, 2 weeks ago

Selected Answer: BE

B= fix every second month issue
E= fix inconsistent connectivity issue.

upvoted 1 times

Ll123123 1 month, 3 weeks ago

Selected Answer: BE

I choose BE. Because DHCP determine which pool to assign either by the incoming interface of DHCP request for directly connected case or if it is a DHCP relay, the discover will carry the incoming address of the relay server receiving this DHCP discover request or an option 82 field, then the DHCP server shall use either field to determine which pool to assign. The default gateway in this config may not be the interface of the DHCP server (connected case) or the giaddr of the relay server which will be automatically excluded.

upvoted 1 times

Muste 4 months, 1 week ago

Selected Answer: BE

sine we don't know if the default-gateway is in this router we have to exclude the default-gateway address from the pool so the correct answer is B&E

upvoted 1 times

[Removed] 5 months ago

Selected Answer: BC

BC,
When configuring dhcp pool the address of the default router is reserved automatically.
The exhibit does not present enough information to infer that the address configured in the pool is that of a standalone DHCP server.

upvoted 1 times

inteldarvid 5 months ago

Selected Answer: BC

Sorry team, in my previous answer, I tried this in my laboratory, and the correct answer is B and C, because only the server ip has to be excluded. The "E" is not correct, because automatically when we configure the default-router that available address is automatically excluded. The "E" is not correct. I tested this in my lab. Correct option is B and C

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: BE

guys thincking, is B and E: Because its necesary exclude ip server, printer, VM, etc. and default gateway (interface router).

upvoted 2 times

> **inteldarvid** 5 months ago
>
> Sorry team, in my previous answer, I tried this in my laboratory, and the correct answer is B and C, because only the server ip has to be excluded. The "E" is not correct, because automatically when we configure the default-router that available address is automatically excluded. The "E" is not correct. I tested this in my lab. Correct option is B and C
>
> upvoted 1 times

tseen 10 months, 2 weeks ago

Selected Answer: CE

From the question, it is not mentioned if this device is a router or a layer 3 switch, also there is no info as per if the gateway is on an interface of this device, hence the gateway can be on any device, and this device is only used as dhcp server. So I will suggest C and E

upvoted 1 times

PimplePooper 12 months ago

Selected Answer: BC

Answer is BC. E is not applicable, as the DHCP pool could still be utilized by other devices and removing the default gateway will cause connectivity issues on those devices.

upvoted 1 times

Huntkey 1 year, 2 months ago

Selected Answer: BE

The question didn't say the DHCP router is the gateway as well. Excluding the gateway IP is a good idea especially the DHCP server is not in the same segment as the client

upvoted 2 times

petr0s 1 year, 9 months ago

Selected Answer: BC

E is wrong, you cannot reserve a default gateway. So BC correct.

upvoted 3 times

bogd 1 year, 10 months ago

Selected Answer: BC

You can only exclude the statically assigned address, you cannot "exclude a default gateway" from a DHCP server.

upvoted 1 times

Carl1999 1 year, 10 months ago

Selected Answer: BE

Set the "ip dhcp excluded-address" including the default gateway and DNS.

upvoted 1 times

> **bogd** 1 year, 10 months ago
>
> You cannot "exclude a default gateway"... Should be BC
>
> upvoted 1 times

Surfside92 2 years, 1 month ago

At first glance the answer looks like B and C.
However I think its B and E.

From the output the default router address 200.30.40.100 is not excluded from the dhcp pool 2. So a host device could pick up that ip address via dhcp.
And the question states the server is not reachable "consistently" ie not at all. So that has probably happened and all devices with ip address from dhcp pool 2 are isolated within the 200.30.40.0/24 network and not reachable.

Note the same issue exists with dhcp pool 1. It has not excluded 200.30.30.100 from the dhcp pool.
You could argue that Cisco DHCP would ping the default-router the address first - get a reply - and mark it as in use and not hand it out. So that

would suggest B and C work as the correct answer.
But that would be very bad prectice - and what if the router address blocks icmp. The Cisco DHCP service would think the address is not in use.

upvoted 2 times

MrThinMints 1 year, 11 months ago

Per cisco doc: "The IP address configured on the router interface is automatically excluded from the DHCP address pool. The DHCP server assumes that all other IP addresses in a DHCP address pool subnet are available for assigning to DHCP clients." So no, you do not have to manually exclude the default gateway, assuming that it is an ip address on the router itself.
https://www.cisco.com/en/US/docs/ios/12_4t/ip_addr/configuration/guide/htdhcpsv.html#:~:text=The%20IP%20address%20configured%20on,for%20assigning%20to%20DHCP%20clients.

upvoted 2 times

tseen 10 months, 2 weeks ago

From the question, it is not mentioned if this device is a router or a layer 3 switch, also there is no info as per if the gateway is on an interface of this device, hence the gateway can be on any device, and this device is only used as dhcp server. So I will suggest B and E

upvoted 1 times

tseen 10 months, 2 weeks ago

I mean C and E not B and E

upvoted 1 times

Networkingguy 1 year, 10 months ago

Great find MrThinMints, so yeah Lock in B and C Eddie

upvoted 1 times

MP_iBGP 2 years, 2 months ago

BE is correct response. IP address of gateway must be also excluded

upvoted 2 times

[Removed] 1 year, 11 months ago

You still have to configure the router to exclude the server address. From the router perspective, if you choose B&E it will still see the server ip being available and try to lease it to requesting devices further creating issues. B,C are correct.

upvoted 1 times

[Removed] 1 year, 11 months ago

LOL disregard comment... I still stay with B,C though. If default gateway is configured on interface that excludes it from the pool.

upvoted 2 times

error_909 2 years, 3 months ago

The given answer is correct

upvoted 1 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 1 times

Question #199 *Topic 1*

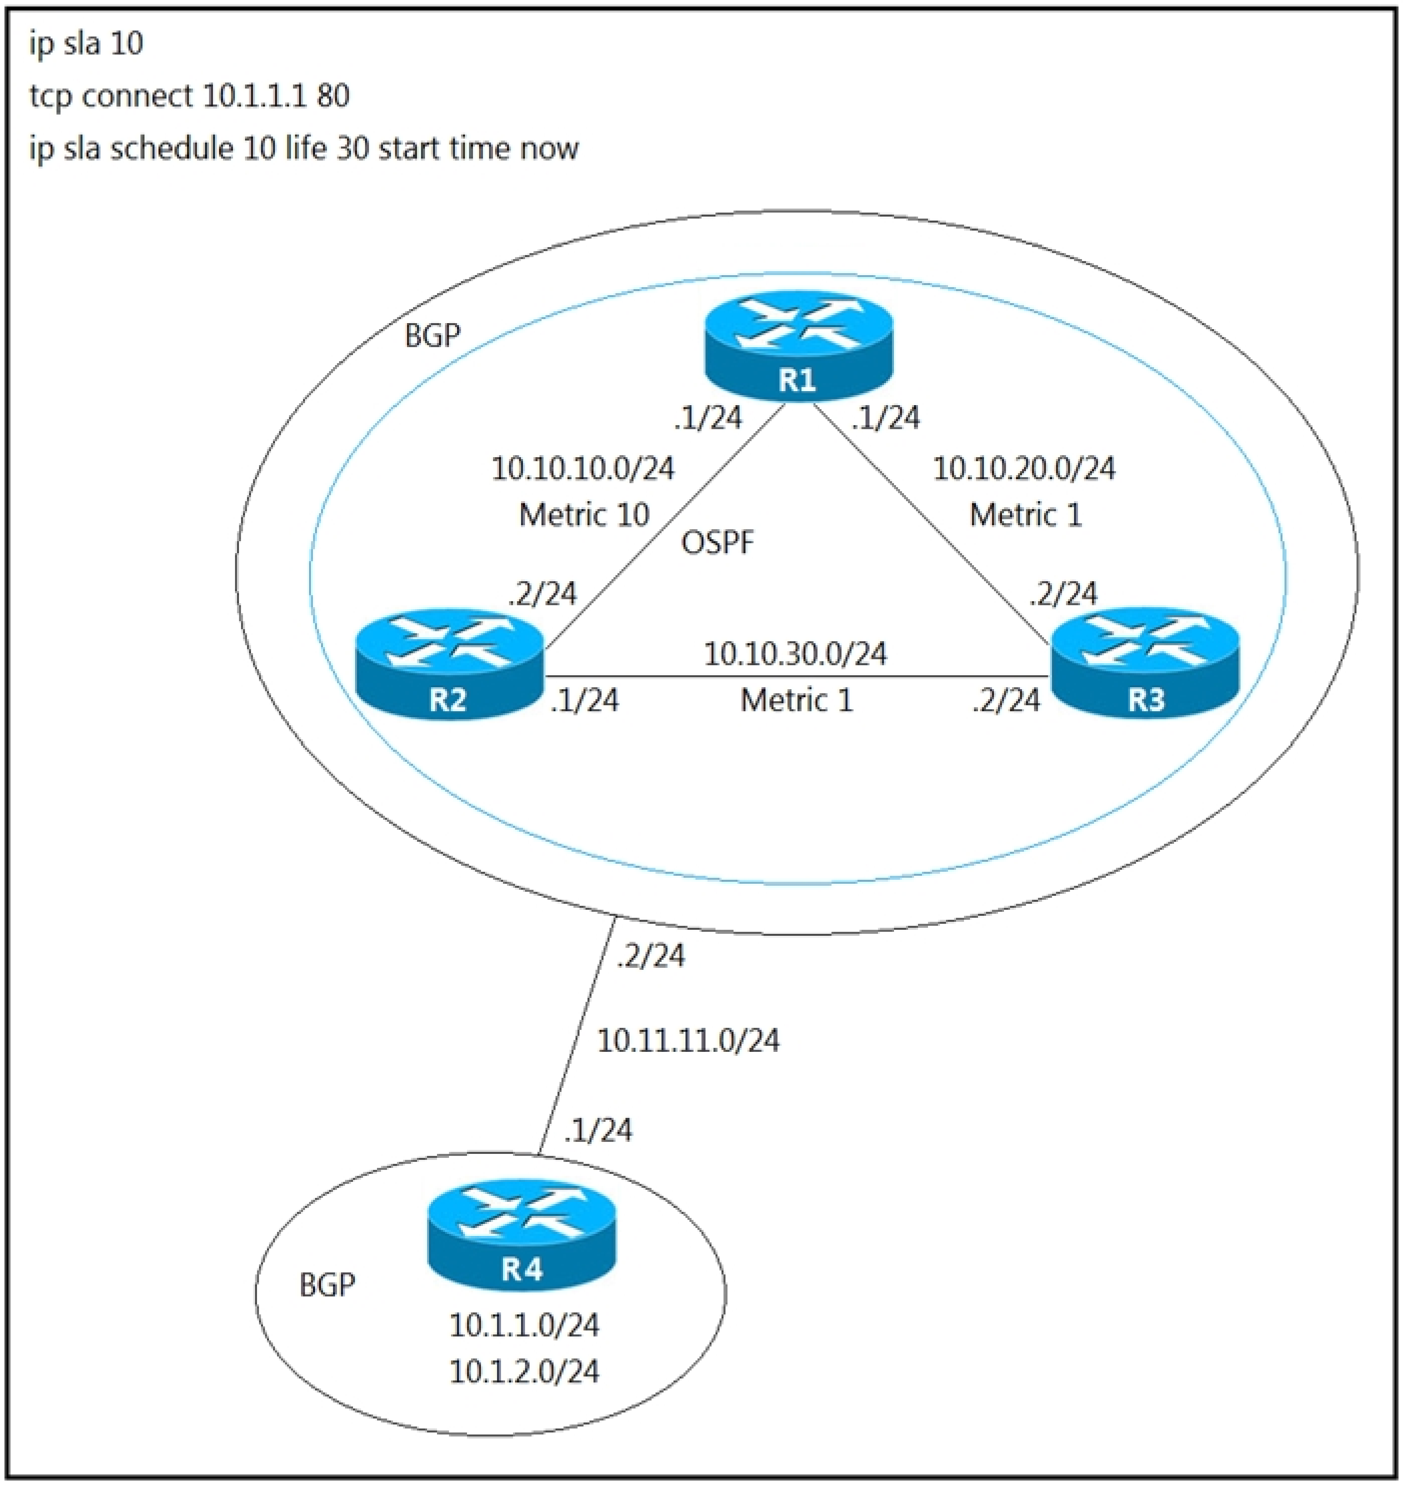

```
ip sla 10
tcp connect 10.1.1.1 80
ip sla schedule 10 life 30 start time now
```

Refer to the exhibit. A user has set up an IP SLA probe to test if a non SLA host web server on IP address 10.1.1.1 accepts HTTP sessions prior to deployment.

The probe is failing.

Which action should the network administrator recommend for the probe to succeed?

A. Re-issue the ip sla schedule command.

B. Add the control disable option to the tcp connect.

C. Modify the ip sla schedule frequency to forever.

D. Add icmp-echo command for the host.

Correct Answer: *A*

Community vote distribution

B (100%)

DaanB Highly Voted 2 years, 8 months ago

According to this Cisco link https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2018/pdf/BRKNMS-3043.pdf, we should disable the Control Protocol with "control disable" keyword (the full command is "tcp connect 10.1.1.1 80 control disable) if the target host is not running IP SLA -> Answer "Add the control disable option to the tcp connect" is correct.

upvoted 22 times

Masashi_O 2 years, 6 months ago

p.28
If the target host is not running IP SLA,disable the Control Protocol (optional).
Default: enabled

upvoted 2 times

5566 Highly Voted 2 years, 3 months ago

Answer is correct: A
ip sla schedule 10 life 30 start-time now
is " start-time" not " start time"

upvoted 7 times

inteldarvid Most Recent 5 months, 2 weeks ago

Selected Answer: B

option B is correct:

https://learningnetwork.cisco.com/s/question/0D53i00000KsuUMCAZ/ipsla-tcpconnect-control-disable

upvoted 1 times

hoins 8 months, 2 weeks ago

Selected Answer: B

The correct answer is B.

upvoted 1 times

WAKIDI 1 year, 5 months ago

reference for "B" : https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/configuration/15-mt/sla-15-mt-book/sla_tcp_conn-0.html#GUID-BF314AA6-8CDD-4941-A4C7-2801CC44D6F1

upvoted 1 times

Eric0_0 1 year, 9 months ago

Selected Answer: B

Correct answer is B. The server is non IP SLA, so control disable is needed when probing the server.

upvoted 2 times

xerex 1 year, 11 months ago

Selected Answer: B

According to this Cisco link, we should disable the Control Protocol with "control disable" keyword (the full command is "tcp connect 10.1.1.1 80 control disable) if the target host is not running IP SLA -> Answer "Add the control disable option to the tcp connect" is correct.

upvoted 3 times

myrmike 1 year, 11 months ago

B seems to be the only logical answer. Both A and C imply that the sla schedule is not running so the probe would not be running and could not fail.

upvoted 1 times

Dirkd0344 2 years ago

The correct answer is B. It allows the operation to perform without configuring a responder on the remote device.

upvoted 1 times

donjime 2 years, 2 months ago

The correct Answer is B, because the HTTP Server it's not running IP SLA so the command Control disable must be enable

upvoted 3 times

error_909 2 years, 3 months ago

Answer is correct: A

upvoted 1 times

examShark 2 years, 4 months ago

The correct answer is B
https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2018/pdf/BRKNMS-3043.pdf

upvoted 4 times

Question #200

Topic 1

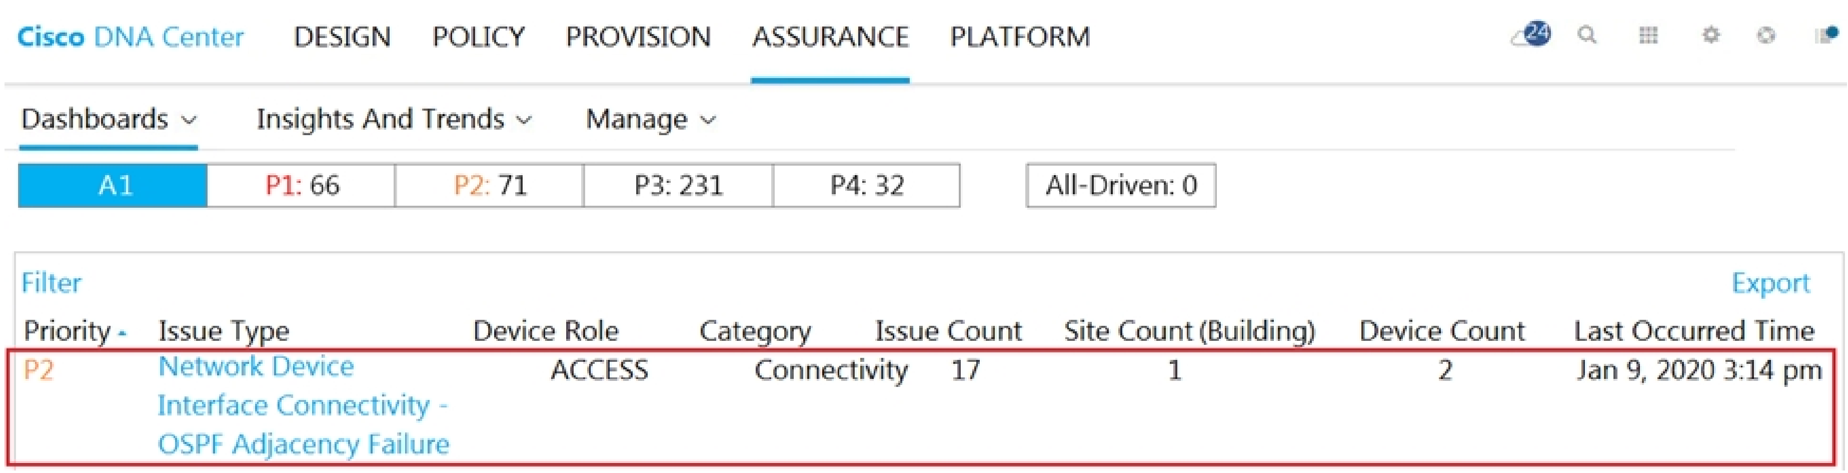

Priority	Issue Type	Device Role	Category	Issue Count	Site Count (Building)	Device Count	Last Occurred Time
P2	Network Device Interface Connectivity - OSPF Adjacency Failure	ACCESS	Connectivity	17	1	2	Jan 9, 2020 3:14 pm

Refer to the exhibit. A network administrator is using the DNA Assurance Dashboard panel to troubleshoot an OSPF adjacency that failed between Edge_NYC

Interface GigabitEthernet1/3 with Neighbor Edge_SNJ. The administrator observes that the neighborship is stuck in the exstart state.

How does the administrator fix this issue?

A. Configure to match the OSPF interface network types on both routers.

B. Configure to match the OSPF interface speed and duplex settings on both routers.

C. Configure to match the OSPF interface MTU settings on both routers.

D. Configure to match the OSPF interface unique IP address and subnet mask on both routers.

Correct Answer: *C*

yeyuno 8 months, 2 weeks ago

Neighbors Stuck in Exstart/Exchange State

The problem occurs most frequently when you attempt to run OSPF between a Cisco router and another vendor router. The problem occurs when the maximum transmission unit (MTU) settings for neighboring router interfaces do not match. If the router with the higher MTU sends a packet larger that the MTU set on the neighboring router, the neighbor router ignores the packet. When this problem occurs, the output of the show ip ospf neighbor command displays output similar to what is shown in this figure.

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct

upvoted 2 times

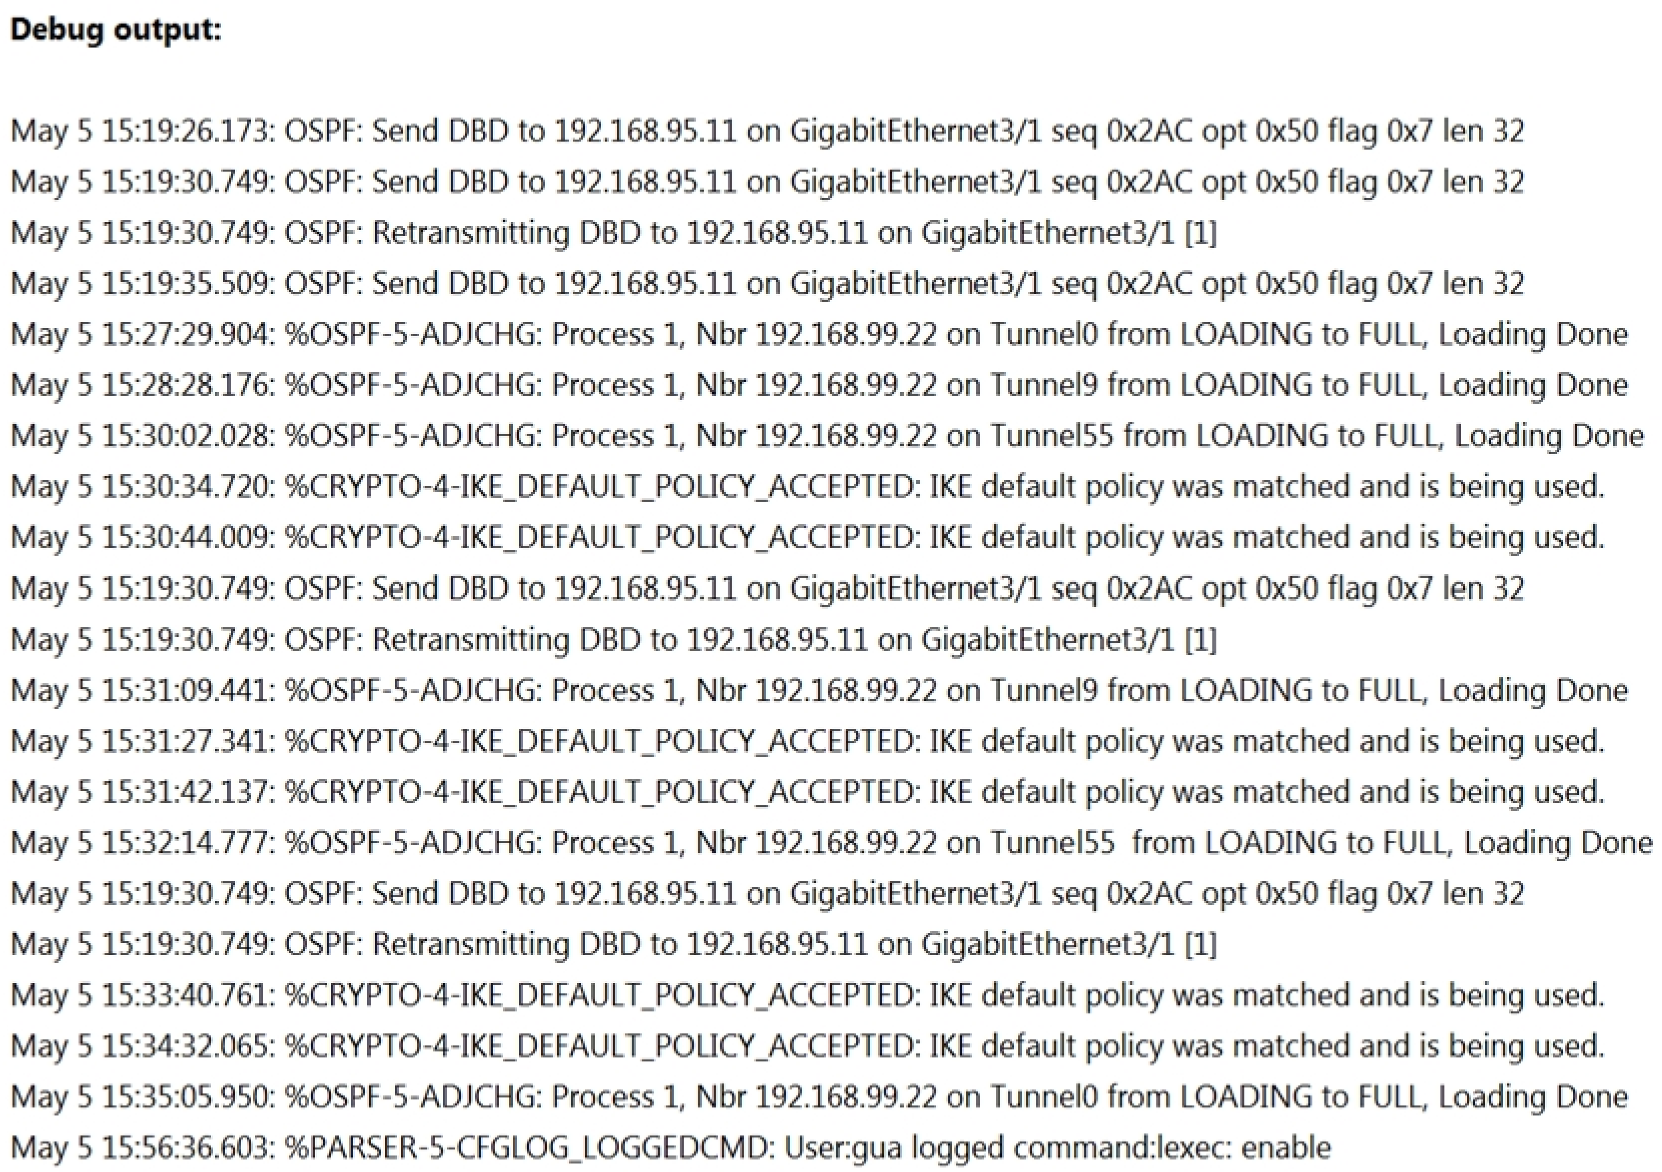

Debug output:

```
May 5 15:19:26.173: OSPF: Send DBD to 192.168.95.11 on GigabitEthernet3/1 seq 0x2AC opt 0x50 flag 0x7 len 32
May 5 15:19:30.749: OSPF: Send DBD to 192.168.95.11 on GigabitEthernet3/1 seq 0x2AC opt 0x50 flag 0x7 len 32
May 5 15:19:30.749: OSPF: Retransmitting DBD to 192.168.95.11 on GigabitEthernet3/1 [1]
May 5 15:19:35.509: OSPF: Send DBD to 192.168.95.11 on GigabitEthernet3/1 seq 0x2AC opt 0x50 flag 0x7 len 32
May 5 15:27:29.904: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.99.22 on Tunnel0 from LOADING to FULL, Loading Done
May 5 15:28:28.176: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.99.22 on Tunnel9 from LOADING to FULL, Loading Done
May 5 15:30:02.028: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.99.22 on Tunnel55 from LOADING to FULL, Loading Done
May 5 15:30:34.720: %CRYPTO-4-IKE_DEFAULT_POLICY_ACCEPTED: IKE default policy was matched and is being used.
May 5 15:30:44.009: %CRYPTO-4-IKE_DEFAULT_POLICY_ACCEPTED: IKE default policy was matched and is being used.
May 5 15:19:30.749: OSPF: Send DBD to 192.168.95.11 on GigabitEthernet3/1 seq 0x2AC opt 0x50 flag 0x7 len 32
May 5 15:19:30.749: OSPF: Retransmitting DBD to 192.168.95.11 on GigabitEthernet3/1 [1]
May 5 15:31:09.441: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.99.22 on Tunnel9 from LOADING to FULL, Loading Done
May 5 15:31:27.341: %CRYPTO-4-IKE_DEFAULT_POLICY_ACCEPTED: IKE default policy was matched and is being used.
May 5 15:31:42.137: %CRYPTO-4-IKE_DEFAULT_POLICY_ACCEPTED: IKE default policy was matched and is being used.
May 5 15:32:14.777: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.99.22 on Tunnel55  from LOADING to FULL, Loading Done
May 5 15:19:30.749: OSPF: Send DBD to 192.168.95.11 on GigabitEthernet3/1 seq 0x2AC opt 0x50 flag 0x7 len 32
May 5 15:19:30.749: OSPF: Retransmitting DBD to 192.168.95.11 on GigabitEthernet3/1 [1]
May 5 15:33:40.761: %CRYPTO-4-IKE_DEFAULT_POLICY_ACCEPTED: IKE default policy was matched and is being used.
May 5 15:34:32.065: %CRYPTO-4-IKE_DEFAULT_POLICY_ACCEPTED: IKE default policy was matched and is being used.
May 5 15:35:05.950: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.99.22 on Tunnel0 from LOADING to FULL, Loading Done
May 5 15:56:36.603: %PARSER-5-CFGLOG_LOGGEDCMD: User:gua logged command:lexec: enable
```

Refer to the exhibit. A network administrator is troubleshooting OSPF adjacency issue by going through the console logs in the router, but due to an overwhelming log messages stream, it is impossible to capture the problem.

Which two commands reduce console log messages to relevant OSPF neighbor problem details so that the issue can be resolved? (Choose two.)

A. debug condition ospf neighbor

B. debug condition interface

C. debug condition session-id ADJCHG

D. debug condition all

Correct Answer: *AB*

Community vote distribution

BC (83%) Other

ridonak230 3 months ago

Selected Answer: BC

B and C are the correct ones !

upvoted 1 times

robi1020 5 months, 2 weeks ago

Selected Answer: BC

Its B and C, Why? B is logical and C (https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/isg/configuration/xe-3s/asr1000/isg-xe-3s-asr1000-book/isg-debug-dcd.pdf)

upvoted 1 times

inteldarvid 5 months ago

genius. Thank you

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

I do not see the "debug condition session-id" command to be related to troubleshooting an OSPF adjacency issue. They probably meant this command:
#debug ip ospf adj

upvoted 2 times

c946f3e 8 months ago

debug condition seems ot be the only valid answer here... except i am missing something

```
site1#debug condition ?
called called number
calling calling
cpl Cisco Provisioning Language debugging
glbp interface group
interface interface
ip IP address
mac-address MAC address
match-list apply the match-list
profile Media Services Profile
standby interface group
username username
vcid VC ID
vrf Virtual Routing and Forwarding
xconnect Xconnect conditional debugging on segment pair
```

upvoted 1 times

yeyuno 8 months ago

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/isg/configuration/xe-3s/asr1000/isg-xe-3s-asr1000-book/isg-debug-dcd.pdf

debug condition session-id

upvoted 1 times

MasterMatt 8 months, 2 weeks ago

Can't find the "debug condition session-id ADJCHG" command. Running a 7200 Software (C7200-ADVENTERPRISEK9-M), Version 15.2(4)M7. Sometimes I'm amazed with these questions. Are we really expected to pass with this ton of hit or miss questions? Not referring to examtopics just Cisco questions are worded poorly and designed bad.

upvoted 1 times

mhd96far 7 months, 3 weeks ago

TOTTALY AGREE , did you pass or not yet

upvoted 1 times

Zizu007 11 months, 2 weeks ago

Selected Answer: B

```
R2#debug condition ?
called called number
calling calling
cpl Cisco Provisioning Language debugging
glbp interface group
interface interface
ip IP address
mac-address MAC address
match-list apply the match-list
profile Media Services Profile
standby interface group
username username
vcid VC ID
vrf Virtual Routing and Forwarding
xconnect Xconnect conditional debugging on segment pair

R2#debug condition
```

upvoted 1 times

Huntkey 1 year, 2 months ago

Selected Answer: BC

```
sw#debug condition os?
% Unrecognized command
sw#debug condition session-id ?
<1-4294967295> Session Number for debug filtering
```

upvoted 1 times

TECH3K3 1 year, 4 months ago

Selected Answer: BC

B & C:
I checked Cisco IOS and IOS-XE and only B & C are valid commands.

upvoted 2 times

JOKERR 1 year, 7 months ago

Selected Answer: CD

The other commands are not valid. Only 2 commands valid are C and D.

upvoted 1 times

JOKERR 1 year, 7 months ago

Apologies. Answer is B and C.

upvoted 2 times

xziomal9 1 year, 7 months ago

Selected Answer: BC

Correct: BC

upvoted 1 times

YaPet 1 year, 10 months ago

Selected Answer: BC

B,C,D are valid commands, A hasn't presented in IOS reference. It seems B,C are best for answer.
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/debug/command/a1/db-a1-cr-book/db-c1.html

upvoted 4 times

JoeyT 6 months, 1 week ago

wrong! "all" only goes with "no debug....."

upvoted 1 times

OakA1 2 years, 2 months ago

There is no such a command debug condition ospf neighbor

upvoted 3 times

beatido 2 years, 3 months ago

The only valid command here is:

debug condition interface

Tried on different IOS versions

upvoted 3 times

Alnet 2 years, 1 month ago

You're right. I haven't been able to find the other commands anywhere. Messed up question.

upvoted 1 times

[Removed] 1 year, 11 months ago

Just verified.. B,C,D are all valid commands (minus ADJCHG on C) so A is definitely eliminated. This question is busted but B&C are the best answers to me..

upvoted 3 times

[Removed] 1 year, 11 months ago

debug condition session-id is a command but the only input was numbers and not words (at least on my current IOS)

upvoted 1 times

error_909 2 years, 3 months ago

Given Answer is incorrect.

It's B&C

upvoted 2 times

examShark 2 years, 4 months ago

The given answer is correct
The session identifier is a unique dynamic number generated internally by the Cisco IOS software and assigned to each session when the session is established. (Not a number)

upvoted 2 times

Question #202 Topic 1

Refer to the exhibit.

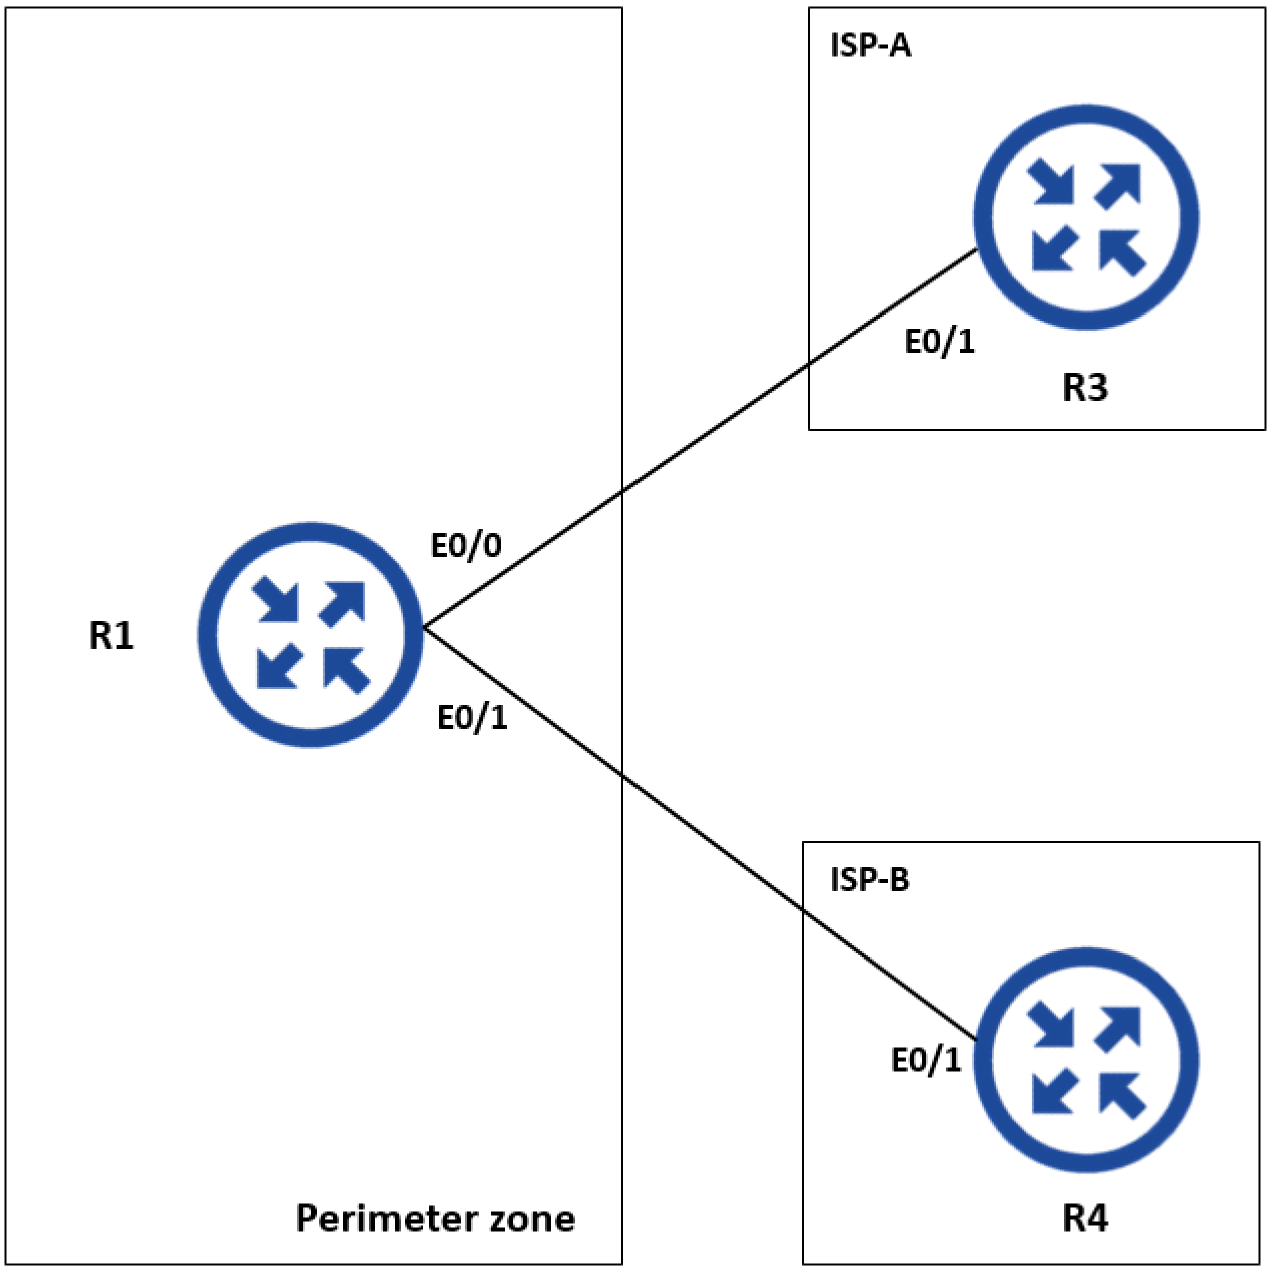

A network is under a cyberattack. A network engineer connected to R1 by SSH and enabled the terminal monitor via SSH session to find the source and destination of the attack. The session was flooded with messages, which made it impossible for the engineer to troubleshoot the issue.

Which command resolves this issue on R1?

A. (config)#terminal no monitor

B. (config)#no terminal monitor

C. #no terminal monitor

D. #terminal no monitor

Correct Answer: *D*

Reference:

https://www.oreilly.com/library/view/cisco-ios-in/0596008694/re826.html

Community vote distribution

D (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

D is correct

Test in my swicth:

```
r-arhouser8-1#terminal no moni
r-arhouser8-1#terminal no monitor
r-arhouser8-1#
```

upvoted 1 times

Lilienen 10 months, 2 weeks ago

Selected Answer: D

Correct answer, tested in lab: #terminal no monitor
It is indeed a non standard order of commands

upvoted 1 times

Nhan 1 year, 5 months ago

So D is correct answer

upvoted 2 times

Nhan 1 year, 5 months ago

Turning Terminal Logging Off

In a classic moment of IOS madness, if you want to stop logging to your terminal:

s1#terminal no monitor

If is was consistent with everything in IOS, you might expect to use:

s1#no terminal monitor †But you would be wrong. This syntax is very old and predates the more standardised IOS conventions.

upvoted 1 times

JingleJangus 1 year, 10 months ago

Selected Answer: D

D is correct. I tested it out despite never using the command haha.

upvoted 3 times

Question #203 *Topic 1*

Refer to the exhibit.

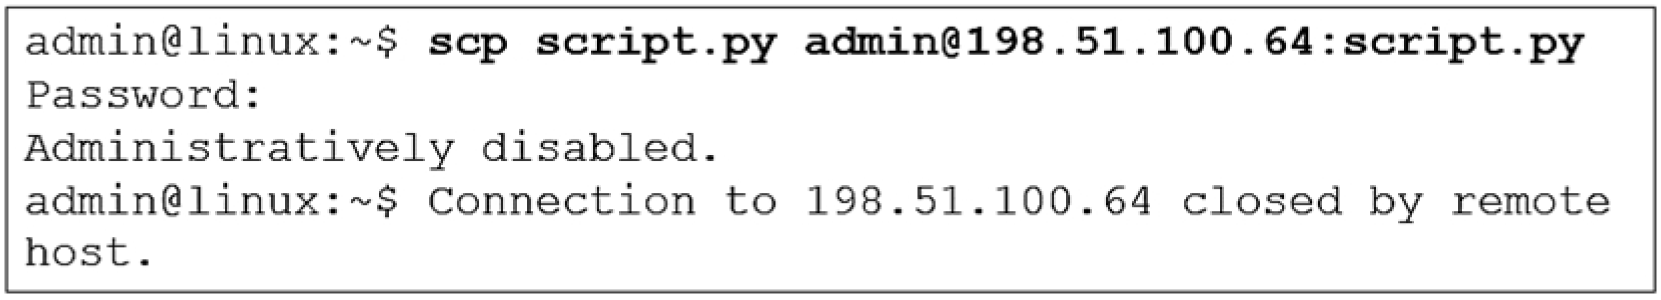

```
admin@linux:~$ scp script.py admin@198.51.100.64:script.py
Password:
Administratively disabled.
admin@linux:~$ Connection to 198.51.100.64 closed by remote
host.
```

A network administrator has developed a Python script on the local Linux machine and is trying to transfer it to the router. However, the transfer fails.

Which action resolves this issue?

A. The Python interpreter must first be enabled with the guestshell enable command.

B. The SSH access must be allowed on the VTY lines using the transport input ssh command.

C. The SSH service must be enabled with the crypto key generate rsa command.

D. The SCP service must be enabled with the ip scp server enable command.

Correct Answer: *D*

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_usr_ssh/configuration/xe-3s/sec-usr-ssh-xe-3s-book/sec-usr-ssh-sec-copy.pdf

Community vote distribution

D (83%) C (17%)

night_wolf_in 1 month, 2 weeks ago

Selected Answer: C

I go with C. We don't need SCP server, the router is a client, and needs to accept the transfer of file. similar to FTP or TFTP, we don't need server to be configured if it is client. SSH is prerequisite for SCP to work.

upvoted 1 times

Pietjeplukgeluk 1 week, 5 days ago

If the Linux server sends a file to the router, the server acts as a client and the router needs to be configured as an SCP server for this to work.

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

D correct:

ip scp server enable
no ip scp server enable

upvoted 1 times

HungarianDish 7 months, 3 weeks ago

Selected Answer: D

https://www.oreilly.com/library/view/cisco-ios-in/0596008694/re451.html

upvoted 3 times

PimplePooper 12 months ago

Selected Answer: D

D is the correct answer

upvoted 1 times

Question #204 *Topic 1*

Refer to the exhibit.

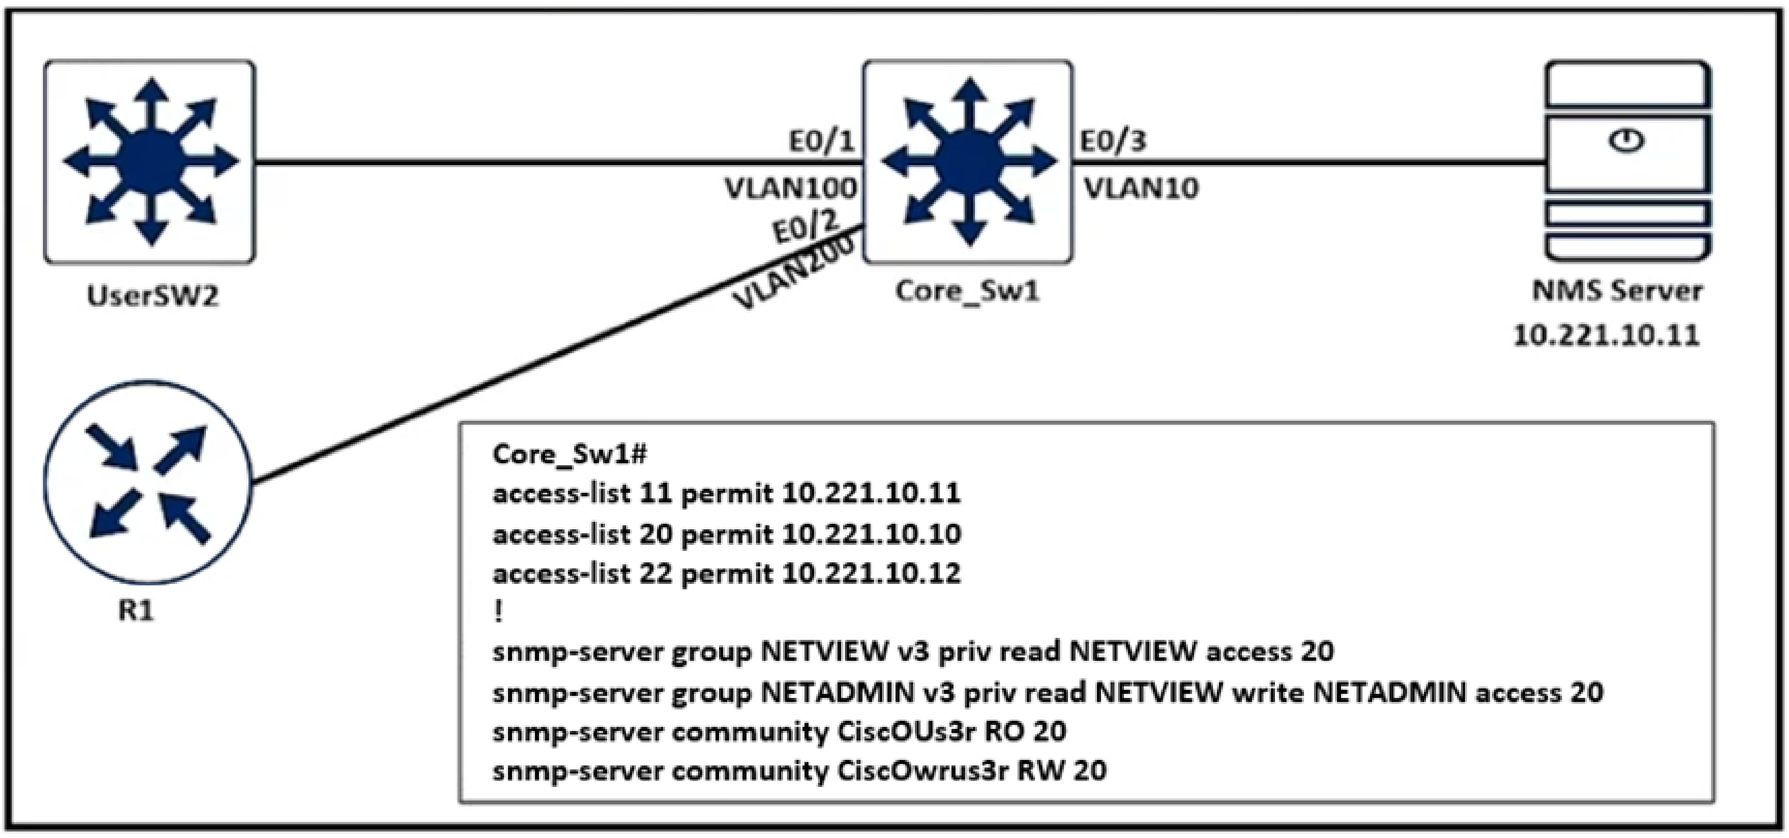

```
Core_Sw1#
access-list 11 permit 10.221.10.11
access-list 20 permit 10.221.10.10
access-list 22 permit 10.221.10.12
!
snmp-server group NETVIEW v3 priv read NETVIEW access 20
snmp-server group NETADMIN v3 priv read NETVIEW write NETADMIN access 20
snmp-server community CiscOUs3r RO 20
snmp-server community CiscOwrus3r RW 20
```

An engineer configured SNMP communities on the Core_Sw1, but the SNMP server cannot obtain information from Core_Sw1.
Which configuration resolves this issue?

A. access-list 20 permit 10.221.10.11

B. snmp-server group NETADMIN v3 priv read NETVIEW write NETADMIN access 22

C. snmp-server group NETVIEW v2c priv read NETVIEW access 20

D. access-list 20 permit 10.221.10.12

Correct Answer: *A*
The SNMP server configuration ties ACL 20 to the list of allowed SNMP servers that can pull data from the switch. The IP address of the NMS server needs to be added to this ACL.

Community vote distribution

A (100%)

inteldarvid 5 months, 2 weeks ago
Selected Answer: A
A is correct. but this question is worng because the rule is duplicate
upvoted 1 times

potato_inet0 7 months, 2 weeks ago
The configuration does not have SNMPv3 users created, which is required, that means if we put the correct ACL on the snmp-group it will not change much, the correct answer is to change the ACL since we do not know if snmpv2c or snmpv3 is used
upvoted 1 times

Dacusai 7 months, 3 weeks ago
Is not a tricky question but it make more sense to put access list 11 on the SNMP configuration that have 2 access lists with the same IP and doing the same thing, no make sense
upvoted 1 times

Question #205 Topic 1

Refer to the exhibit.

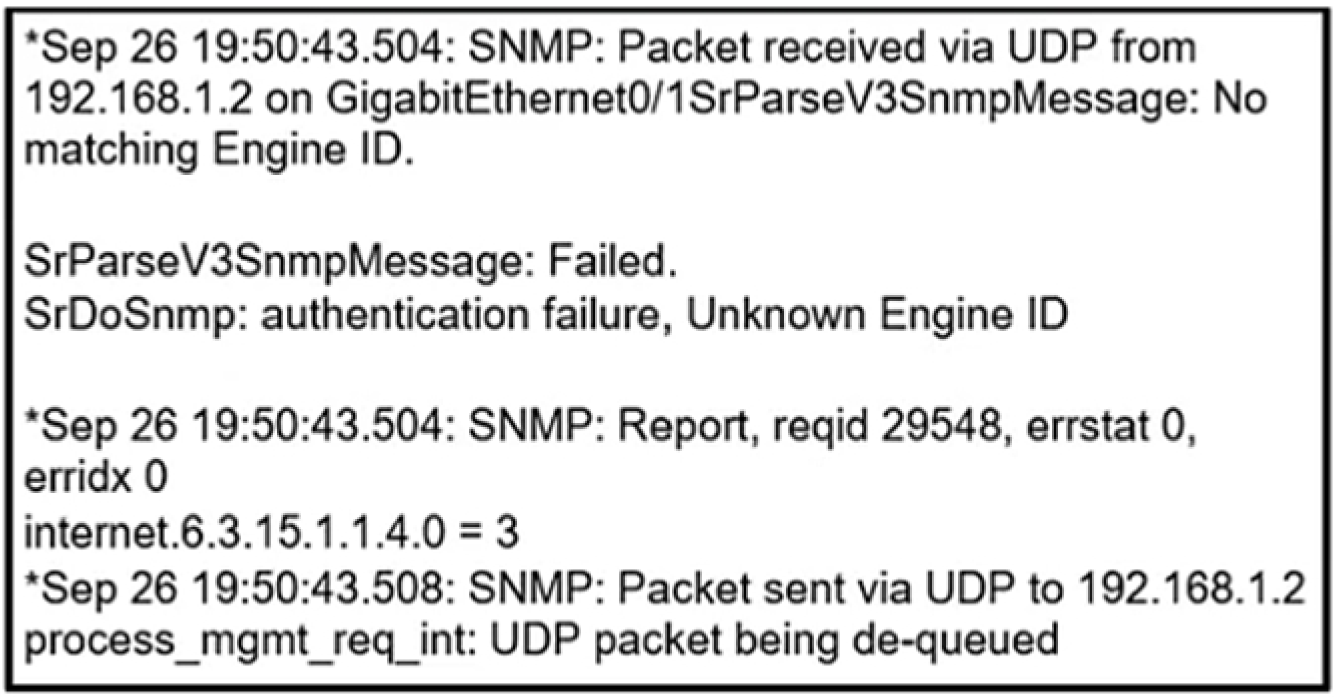

```
*Sep 26 19:50:43.504: SNMP: Packet received via UDP from
192.168.1.2 on GigabitEthernet0/1SrParseV3SnmpMessage: No
matching Engine ID.

SrParseV3SnmpMessage: Failed.
SrDoSnmp: authentication failure, Unknown Engine ID

*Sep 26 19:50:43.504: SNMP: Report, reqid 29548, errstat 0,
erridx 0
internet.6.3.15.1.1.4.0 = 3
*Sep 26 19:50:43.508: SNMP: Packet sent via UDP to 192.168.1.2
process_mgmt_req_int: UDP packet being de-queued
```

Which two commands provide the administrator with the information needed to resolve the issue? (Choose two.)

A. debug snmpv3 engine-id

B. show snmp user

C. debug snmp packet

D. debug snmp engine-id

E. show snmpv3 user

Correct Answer: *BC*

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/command/snmp-xe-3se-3850-cr-book/snmp-xe-3se-3850-cr-book_chapter_0110.html

Community vote distribution

BC (100%)

HungarianDish Highly Voted 7 months, 3 weeks ago

Selected Answer: BC

Example taken from here:
https://community.cisco.com/t5/network-management/snmpv3-not-working/td-p/2934301

upvoted 6 times

ZamanR Most Recent 5 days, 8 hours ago

BC is correct

upvoted 1 times

Refer to the exhibit.

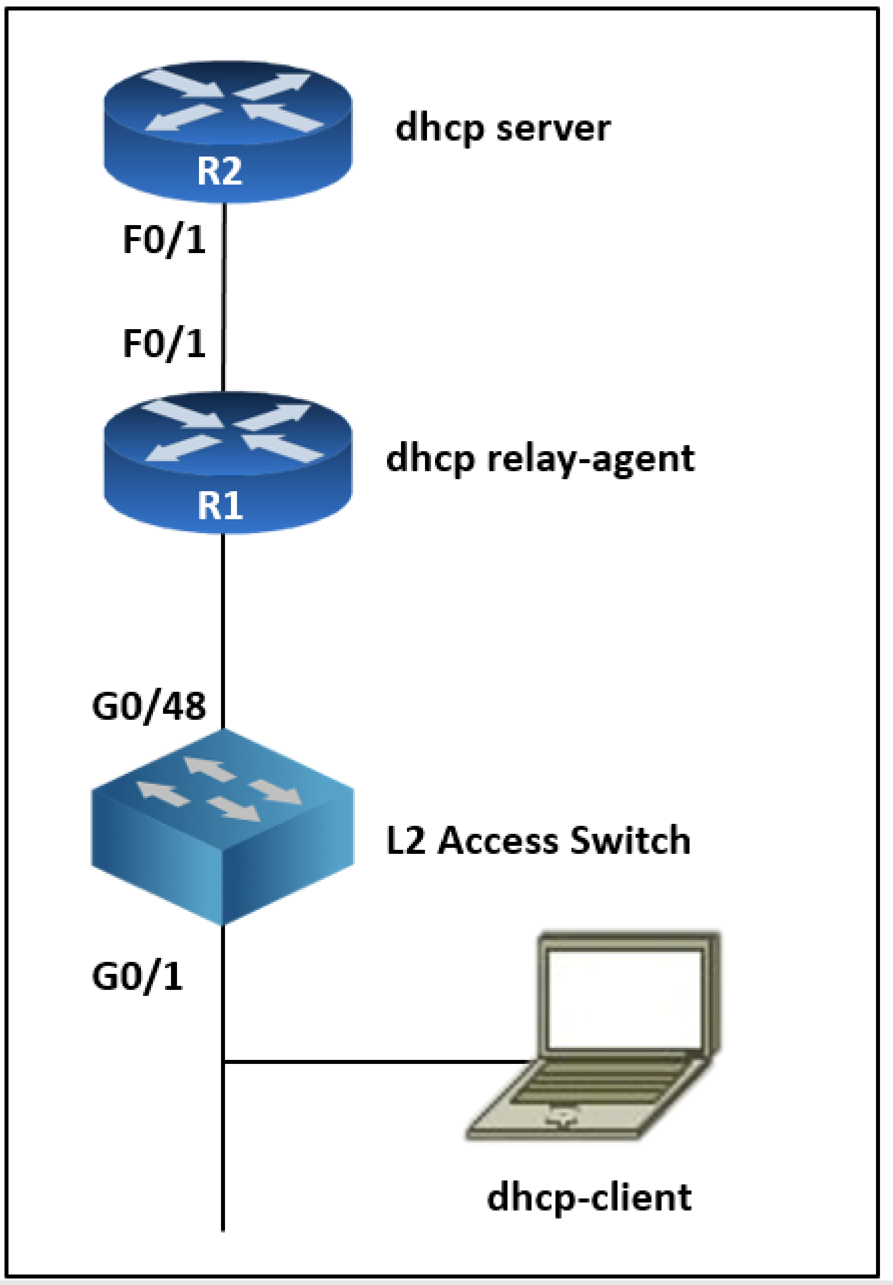

The network administrator can see the DHCP discovery packet in R1, but R2 is not replying to the DHCP request. The R1 related interface is configured with the
DHCP helper address. If the PC is directly connected to the Fa0/1 interface on R2, the DHCP server assigns as IP address from the DHCP pool to the PC.
Which two commands resolve this issue? (Choose two.)

A. service dhcp-relay command on R1

B. ip dhcp relay information enable command on R1

C. ip dhcp option 82 command on R2

D. service dhcp command on R1

E. ip dhcp relay information trust-all command on R2

Correct Answer: *CD*

Community vote distribution

DE (83%) BE (17%)

⊟ tyh391 Highly Voted 1 year, 11 months ago

Answer is D and E

https://community.cisco.com/t5/switching/cisco-router-configured-as-a-dhcp-server-not-replying-to-quot/td-p/3206932

1.- The relay agent was configured by default with "no service dhcp". This caused the relayed packets to come from 0.0.0.0 rather than 10.2.1.1

2.- The DHCP server needs to be configured with "ip dhcp relay information trust-all" so it processes relayed packets with no Giaddr field

upvoted 12 times

⊟ inteldarvid Most Recent 5 months, 2 weeks ago

Selected Answer: DE

Option correct D and E
https://www.cisco.com/en/US/docs/ios/12_4t/ip_addr/configuration/guide/htdhcpre.html#wp1027171
upvoted 1 times

Xerath 9 months, 4 weeks ago

Selected Answer: DE

D & E for sure, references:
ip dhcp relay information trust-all : This command is useful if there is a switch in between the client and the relay agent that may insert option 82. Use this command to ensure that these packets do not get dropped.
https://www.cisco.com/en/US/docs/ios/12_4t/ip_addr/configuration/guide/htdhcpre.html

Prerequisites for Configuring the Cisco IOS DHCP Relay Agent:
The Cisco IOS DHCP server and relay agent are enabled by default. You can verify whether they have been disabled by checking your configuration file. If they have been disabled, the no service dhcp command will appear in the configuration file. Use the service dhcp command to reenable the functionality if necessary.
https://content.cisco.com/chapter.sjs?uri=/searchable/chapter/content/en/us/td/docs/ios-xml/ios/ipaddr_dhcp/configuration/15-sy/dhcp-15-sy-book/configuring_cisco_ios_dhcp_relay_agent.html.xml
upvoted 3 times

TECH3K3 1 year, 5 months ago

Labbed it and none of the answers worked
upvoted 2 times

loklok 1 year, 6 months ago

I agree with DE
upvoted 1 times

cyrus777 1 year, 8 months ago

D and E are correct
A doesn't exit
B doesn't exit
C doesn't exit

```
R2(config)#service dhcp-relay ?
% Unrecognized command
R2(config)#service dhcp-relay

R2(config)#ip dhcp option ?
% Unrecognized command
R2(config)#ip dhcp option
R2(config)#ip dhcp relay information ?
check Validate relay information in BOOTREPLY
option Insert relay information in BOOTREQUEST
policy Define reforwarding policy
trust-all Received DHCP packets may contain relay info option with zero
giaddr

R2(config)#ip dhcp relay information

R2(config)#service dhcp
R2(config)#
```

upvoted 1 times

Eric0_0 1 year, 9 months ago

Selected Answer: DE

Answer is D and E
upvoted 2 times

bogd 1 year, 10 months ago

Selected Answer: DE

See

https://community.cisco.com/t5/switching/cisco-router-configured-as-a-dhcp-server-not-replying-to-quot/td-p/3206932
upvoted 4 times

Carl1999 1 year, 10 months ago

there in no ip dhcp relay information enable command.
#ip dhcp relay information trusted ??

```
(config-if)#ip dhcp relay information ?
check-reply Validate relay information in BOOTREPLY
option DHCP relay information option
option-insert Insert relay information in BOOTREQUEST
policy-action Define reforwarding policy
trusted Received DHCP packet may contain relay info option with zero
giaddr

(config-if)#ip dhcp relay information trusted
```

upvoted 2 times

Carl1999 1 year, 10 months ago

Answer is D and E.
B command does not exist.
It is considered that snooping is set for the l2 switch, so E is required on R2.

upvoted 2 times

[Removed] 1 year, 10 months ago

Selected Answer: BE

B & E..
A - isn't a command C - isn't a real command and depends on the IOS if it supports option 82. D - is wrong, that command enables the DHCP server and its on R1 which is the relay agent.

B adds the DHCP relay agent information option which is (option 82) which is additional info about the relay agent. E allows for requests to process that has a zero for the giaddress.

upvoted 2 times

[Removed] 1 year, 10 months ago

If the relay agent inserts option 82 but does not set the giaddr field in the DHCP packet, the DHCP server interface must be configured as a trusted interface by using the ip dhcp relay information trusted command. This configuration prevents the server from dropping the DHCP message. More info on why its B & E

upvoted 2 times

Question #207 Topic 1

A network administrator performed a Compact Flash Memory upgrade on a Cisco Catalyst 6509 Switch. Everything is functioning normally except SNMP, which was configured to monitor the bandwidth of key interfaces but the interface indexes are changed.
Which global configuration resolves the issue?

A. snmp-server ifindex persist

B. snmp-server ifindex permanent

C. snmp ifindex persist

D. snmp ifindex permanent

Correct Answer: *A*
Reference:
https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/15S/configuration/guide/7600_15_0s_book/ifindx.pdf

Community vote distribution

A (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: A

option A is correct:
https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/15S/configuration/guide/7600_15_0s_book/ifindx.html#:~:text=SNMP%20IfIndex%20Persistence-,Understanding%20SNMP%20IfIndex%20Persistence,a%20physical%20or%20logical%20interface.

upvoted 1 times

NoUserName1234 1 year, 2 months ago

Selected Answer: A

Must be A look a cisco toppic.
https://community.cisco.com/t5/network-management/snmp-errors/td-p/672696

upvoted 1 times

Huntkey 1 year, 2 months ago

Selected Answer: A

```
sw(config)#snmp-server ifindex persist ?
<cr> <cr>
```

upvoted 1 times

Nhan 1 year, 6 months ago

A is correct answer
here is thew sample of real world solarwind snmp server deployment scripts
SOLARWINDS NETFLOW SCRIPTS

```
! from the device’s global configuration mode
! call the int up
int g0/0
! capture inbound traffic
ip flow ingress
! capture outbound traffic
ip flow egress
exit
! definite ip flow export source int
ip flow-export source g0/0
! definite ip flow version
ip flow-export version 5
! definite ip flow export destination
ip flow-export destination 10.10.10.150 2055
! set the flow time out
ip flow-cache timeout active 1
ip flow-cache timeout inactive 15
snmp-server ifindex persist
```

upvoted 2 times

Macferson 1 year, 6 months ago

In the following example, SNMP ifIndex persistence is enabled for Ethernet interface 3/1 only:

```
router(config)# interface ethernet 3/1
router(config-if)# snmp ifindex persist
router(config-if)# exit
```

upvoted 1 times

⊟ **cyrus777** 1 year, 8 months ago

Selected Answer: A

```
R2(config)#snmp-server ifindex persist
R2(config)# snmp-server ifindex permanent
^
% Invalid input detected at '^' marker.

R2(config)#snmp ifindex permanent
^
% Invalid input detected at '^' marker.

R2(config)#
```

upvoted 1 times

⊟ **cyrus777** 1 year, 8 months ago

```
R2(config)# snmp ifindex persist
^
% Invalid input detected at '^' marker.

R2(config)#
```

upvoted 1 times

⊟ **PoopShoot** 1 year, 9 months ago

Actually, I believe the answer is C. The KEY here in the question is KEY interfaces. The snmp-server ifindex persist enables for ALL interfaces.

While snmp ifindex persist is a per interface application.

Link:https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/15S/configuration/guide/7600_15_0s_book/ifindx.html#:~:text=The%20SNMP%20ifIndex%20persistence%20feature%20provides%20an%20interface%20index%20(ifIndex,a%20physical%20or%20logical%20interface.

upvoted 1 times

⊟ **ookr** 1 year, 8 months ago

Yes, but it also says "Which global configuration resolves the issue" and the key here is global.
To me it's A, as it's a global config and enables to all interfaces (and that includes the key interfaces)
But who knows. For some reason Cisco keeps not being clear with the wording in the question. No idea why as this is supposed to be a technical exam rather than en English one.

upvoted 1 times

⊟ **Eric0_0** 1 year, 9 months ago

Selected Answer: A

Correct answer is A. Period.

upvoted 2 times

⊟ **WesleyD** 1 year, 10 months ago

I tested the command at a Cisco 6500, Router(config)# snmp-server ifindex persist is correct

When I give a snmp ?, I don't get the choise for "ifindex"

upvoted 3 times

⊟ **Carl1999** 1 year, 10 months ago

Selected Answer: A

Router(config)# snmp-server ifindex persist

https://www.cisco.com/c/en/us/support/docs/ip/simple-network-management-protocol-snmp/28420-ifIndex-Persistence.html

upvoted 2 times

⊟ **Girmiti** 1 year, 11 months ago

Selected Answer: A

A is correct

upvoted 2 times

⊟ **kkkki** 1 year, 11 months ago

Selected Answer: A

from cisco press

upvoted 1 times

⊟ **geek1992** 1 year, 11 months ago

Is correct ?

upvoted 1 times

⊟ **nial** 1 year, 12 months ago

Correct Answer: snmp-server ifindex persist
To globally enable SNMP ifIndex persistence, perform this task:
router(config)# snmp-server ifindex persist
To enable SNMP ifIndex persistence only on a specific interface, perform this task:
Router(config-if)# snmp ifindex persist

upvoted 3 times

Question #208 *Topic 1*

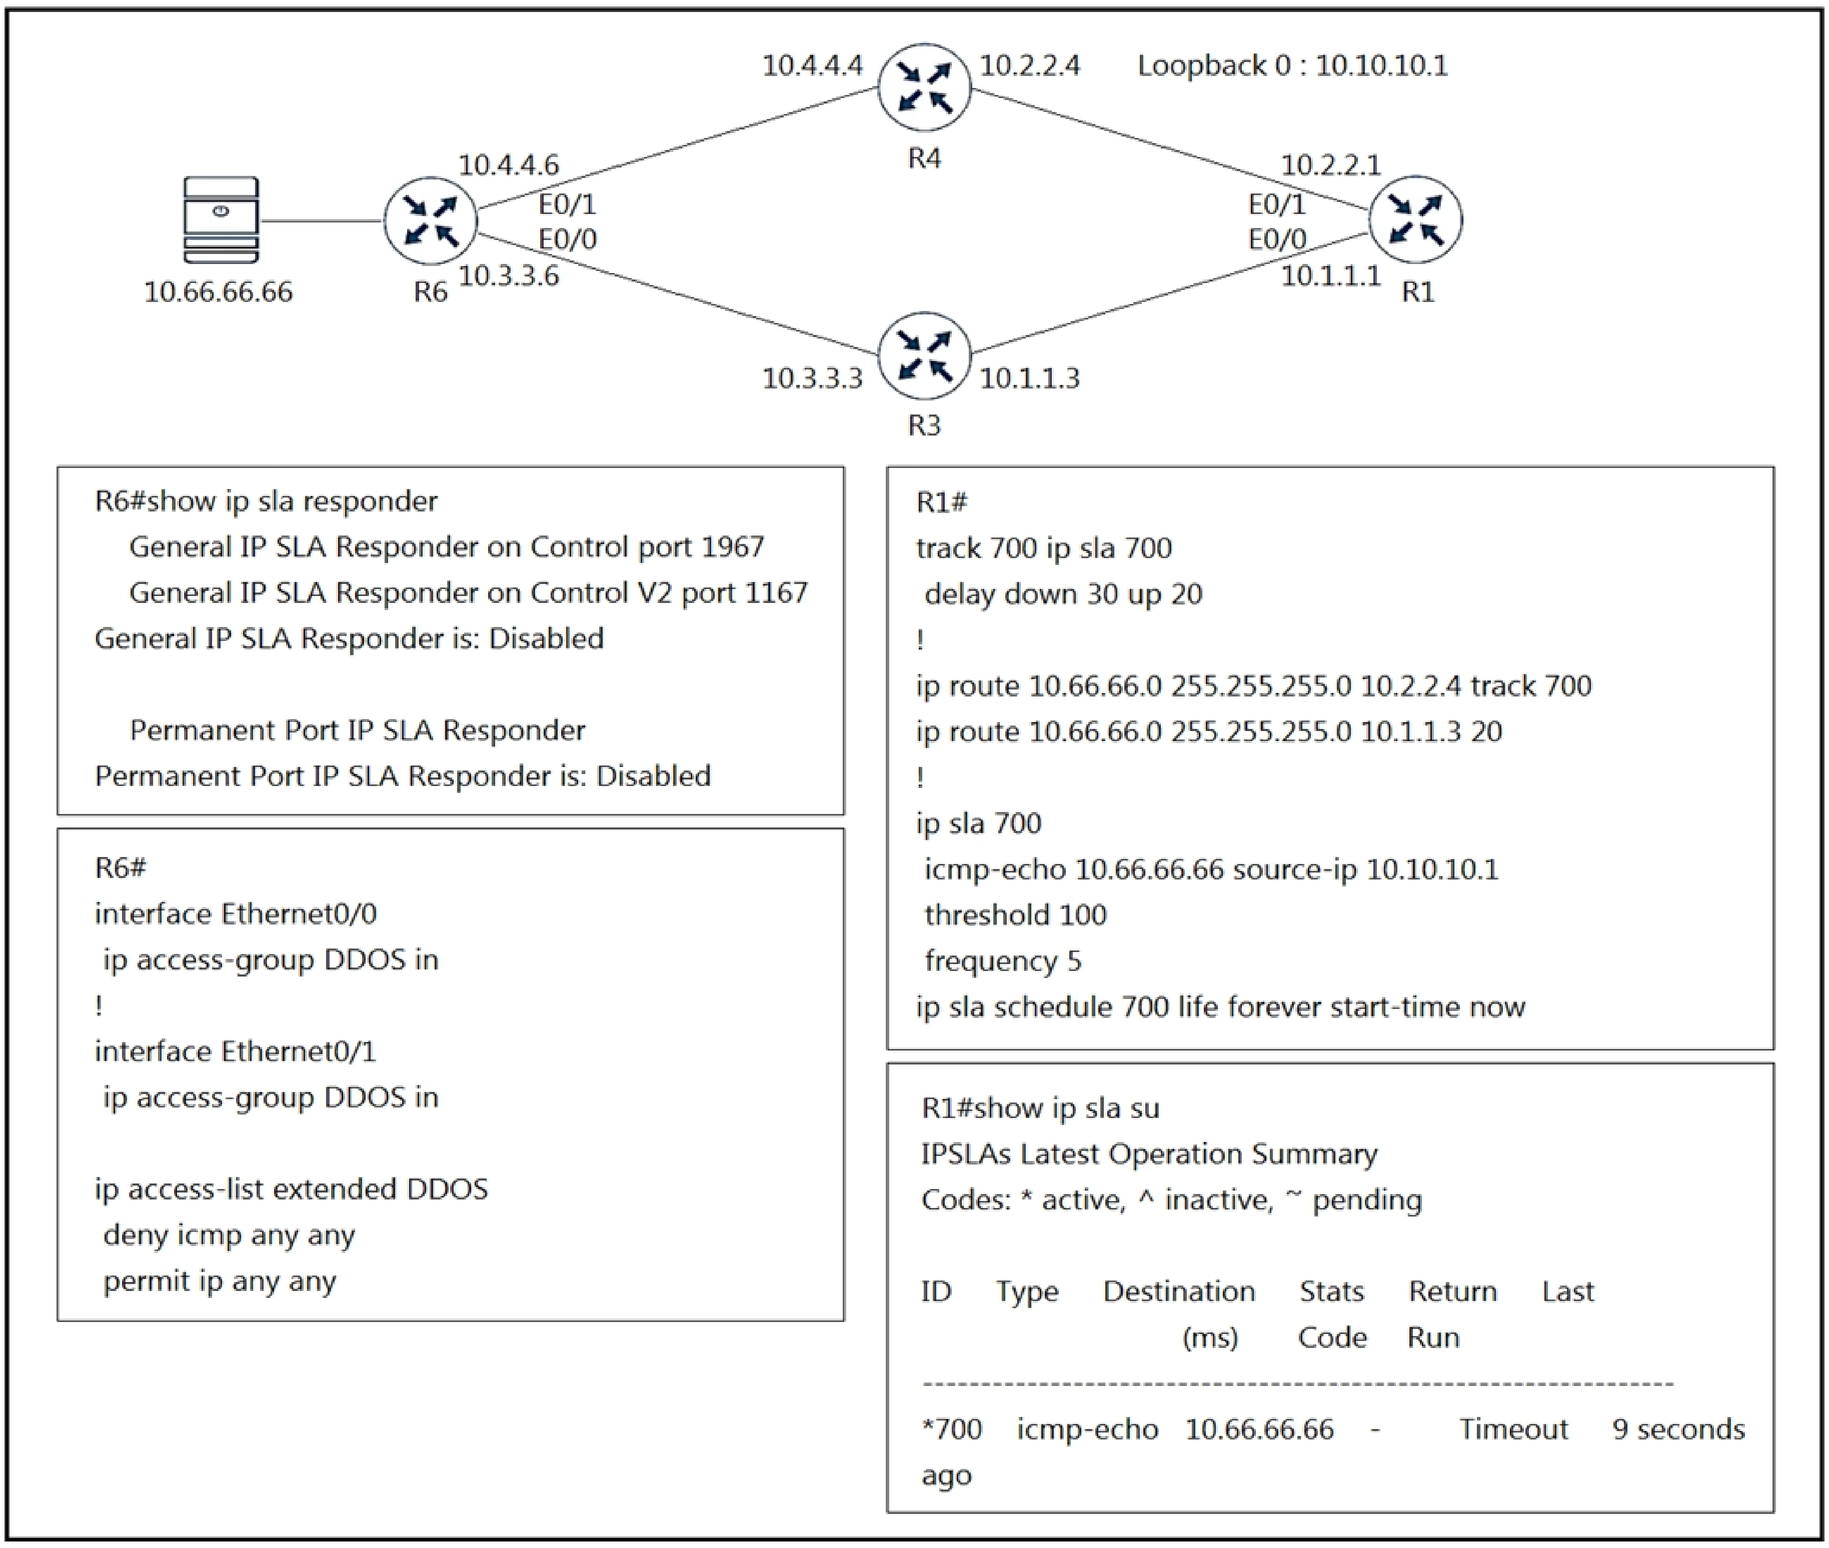

```
R6#show ip sla responder
   General IP SLA Responder on Control port 1967
   General IP SLA Responder on Control V2 port 1167
General IP SLA Responder is: Disabled

   Permanent Port IP SLA Responder
Permanent Port IP SLA Responder is: Disabled
```

```
R6#
interface Ethernet0/0
 ip access-group DDOS in
!
interface Ethernet0/1
 ip access-group DDOS in

ip access-list extended DDOS
 deny icmp any any
 permit ip any any
```

```
R1#
track 700 ip sla 700
 delay down 30 up 20
!
ip route 10.66.66.0 255.255.255.0 10.2.2.4 track 700
ip route 10.66.66.0 255.255.255.0 10.1.1.3 20
!
ip sla 700
 icmp-echo 10.66.66.66 source-ip 10.10.10.1
 threshold 100
 frequency 5
ip sla schedule 700 life forever start-time now
```

```
R1#show ip sla su
IPSLAs Latest Operation Summary
Codes: * active, ^ inactive, ~ pending

ID    Type       Destination    Stats    Return    Last
                       (ms)     Code     Run
-----------------------------------------------------------------
*700  icmp-echo  10.66.66.66    -        Timeout   9 seconds
ago
```

Refer to the exhibit. R1 is configured with IP SLA to check the availability of the server behind R6 but it kept failing.
Which configuration resolves the issue?

A. R6(config)#ip sla responder udp-echo ip address 10.10.10.1 port 5000

B. R6(config)#ip access-list extended DDOS R6(config-ext-nacl)#5 permit icmp host 10.10.10.1 host 10.66.66.66

C. R6(config)#ip sla responder

D. R6(config)#ip access-list extended DDOS R6(config-ext-nacl)#5 permit icmp host 10.66.66.66 host 10.10.10.1

Correct Answer: *B*

Community vote distribution

B (84%) D (16%)

VergilP Highly Voted 1 year, 1 month ago

Selected Answer: B

ACL in R6 E0/0 and E0/1 inbond direction ...
please look the picture carefully .....
source is 10.10.10.1 destination is the 10.66.66.66
B is correct

upvoted 5 times

Huntkey Highly Voted 1 year, 2 months ago

Selected Answer: B

Am I missing something here? R1 uses SLA to send ICMP to the server. The source is 10.10.10.1 and the destination is the 10.66.66.66. I think the ACL in B would perfectly allow it

upvoted 5 times

ZamanR Most Recent 6 days, 7 hours ago

Answer B

In this IP SLA tracking, we don’t need a IP SLA Responder so the command “ip sla responder” on R6 isnot necessary.

We also notice that the ACL is blocking ICMP packets on both interfaces E0/0 & E0/1 of R6 so we need

to allow ICMP from source 10.10.10.1 to destination 10.66.66.66

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

option B is correct team.

upvoted 2 times

ntdevera 1 year, 3 months ago

Selected Answer: D

D, Acl in inwards. Source is the snmp server in that direction.

upvoted 2 times

> **Lilienen** 10 months ago
>
> D is wrong, because ACL is applied to R6, not R1. Review the exhibit properly!
>
> upvoted 1 times

> **quyle** 1 year, 2 months ago
>
> correct, acl in -> source is the snmp server
>
> upvoted 1 times

TECH3K3 1 year, 5 months ago

How is B correct:
The ACL is inbound for both interfaces on R1. So that's server towards R1.
So I would be going for D

upvoted 1 times

> **ericxw** 11 months, 3 weeks ago
>
> do you mean both interfaces on R6?
>
> upvoted 1 times

Nhan 1 year, 5 months ago

B is correct answer, the statement shows that icmp deny deny its will go sequence 10 there for you can set a new statement permit icmp with sequence 5 to allow the traffic because the ACL is being processed by sequency

upvoted 2 times

WAKIDI 1 year, 5 months ago

Selected Answer: B

the "ip sla icmp-echo" in R1 doesn't require an "ip sla responder" in the destination (R6). so A & C wouldn't be appropriate, right ? for "D" , the ACL source and dest addr need to be swapped.

upvoted 4 times

pompedom 1 year, 6 months ago

Selected Answer: D

I think it's D because the acl is configured inward. , traffic going out will not be blocked.

upvoted 1 times

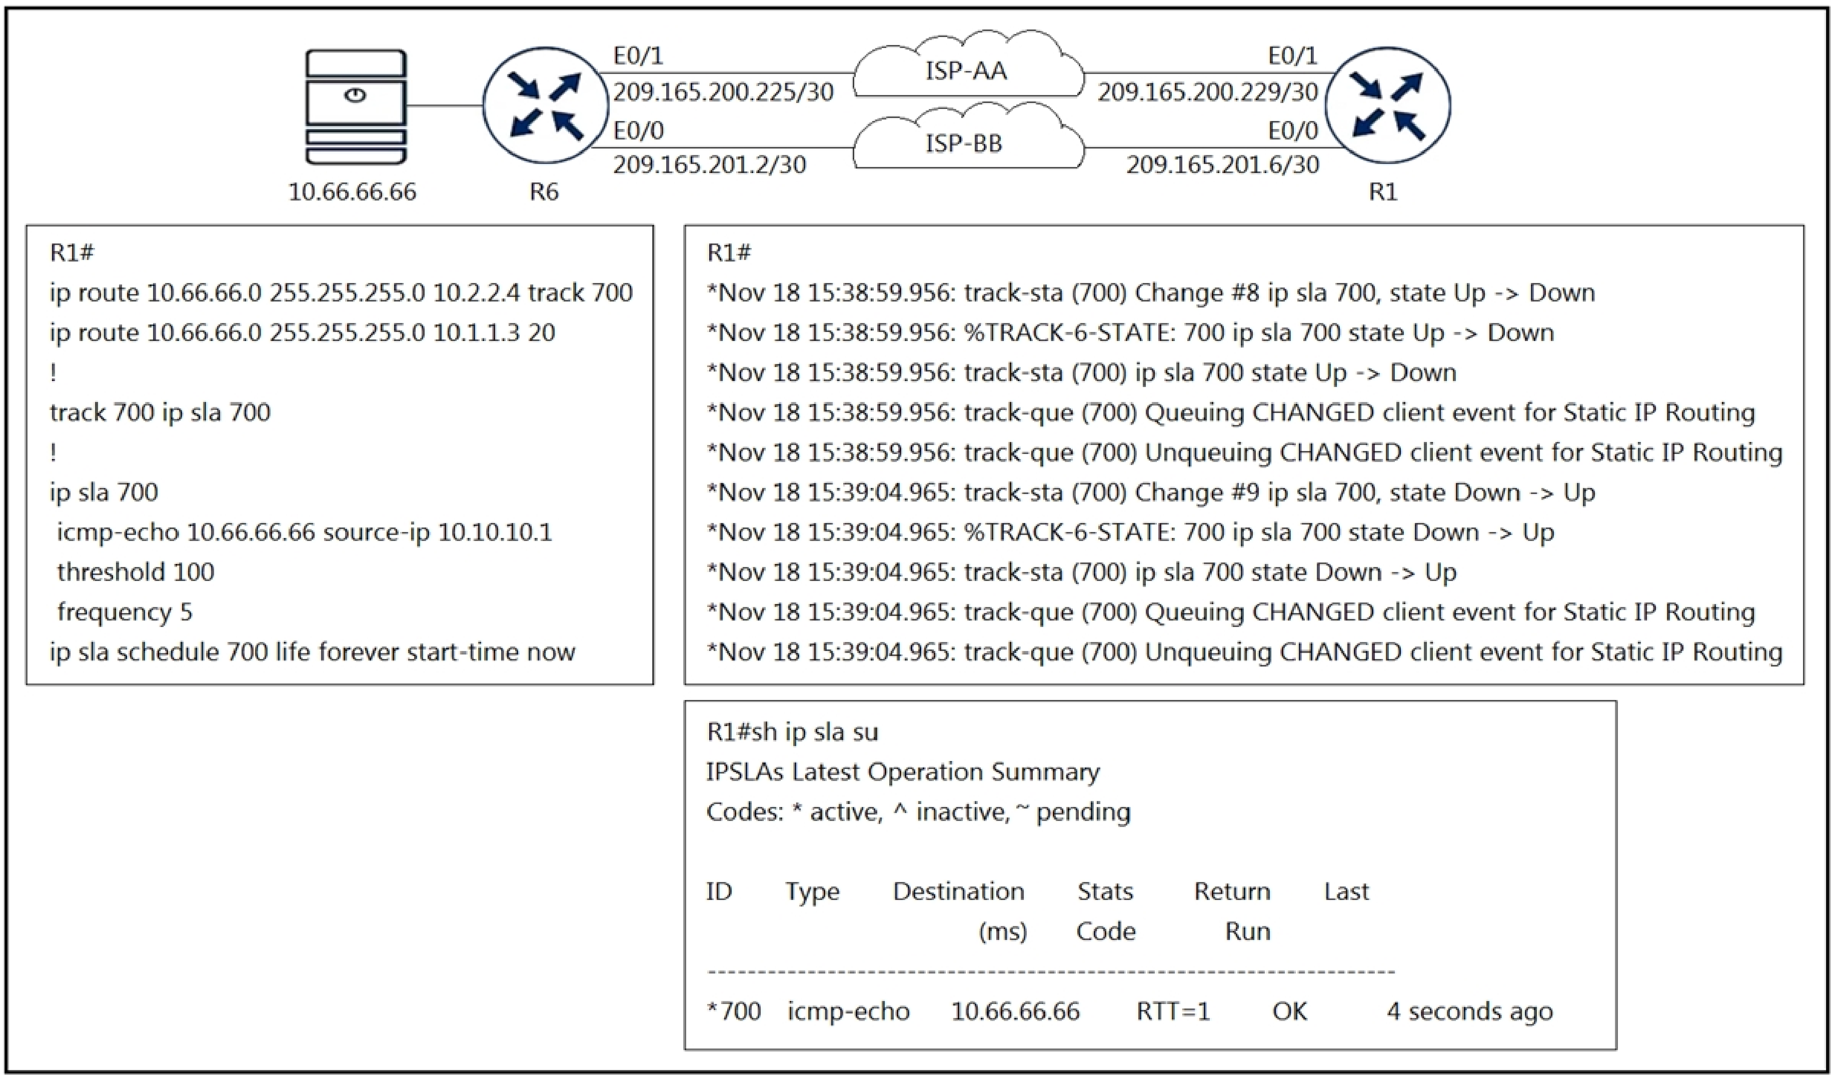

```
R1#
ip route 10.66.66.0 255.255.255.0 10.2.2.4 track 700
ip route 10.66.66.0 255.255.255.0 10.1.1.3 20
!
track 700 ip sla 700
!
ip sla 700
 icmp-echo 10.66.66.66 source-ip 10.10.10.1
 threshold 100
 frequency 5
ip sla schedule 700 life forever start-time now
```

```
R1#
*Nov 18 15:38:59.956: track-sta (700) Change #8 ip sla 700, state Up -> Down
*Nov 18 15:38:59.956: %TRACK-6-STATE: 700 ip sla 700 state Up -> Down
*Nov 18 15:38:59.956: track-sta (700) ip sla 700 state Up -> Down
*Nov 18 15:38:59.956: track-que (700) Queuing CHANGED client event for Static IP Routing
*Nov 18 15:38:59.956: track-que (700) Unqueuing CHANGED client event for Static IP Routing
*Nov 18 15:39:04.965: track-sta (700) Change #9 ip sla 700, state Down -> Up
*Nov 18 15:39:04.965: %TRACK-6-STATE: 700 ip sla 700 state Down -> Up
*Nov 18 15:39:04.965: track-sta (700) ip sla 700 state Down -> Up
*Nov 18 15:39:04.965: track-que (700) Queuing CHANGED client event for Static IP Routing
*Nov 18 15:39:04.965: track-que (700) Unqueuing CHANGED client event for Static IP Routing
```

```
R1#sh ip sla su
IPSLAs Latest Operation Summary
Codes: * active, ^ inactive, ~ pending

ID    Type       Destination     Stats      Return     Last
                         (ms)    Code       Run
-----------------------------------------------------------------------
*700  icmp-echo     10.66.66.66     RTT=1     OK        4 seconds ago
```

Refer to the exhibit. An engineer configured IP SLA on R1 to avoid the ISP link flapping problem, but it is not working as designed. IP SLA should wait 30 seconds before switching traffic to a secondary connection and then revert to the primary link after waiting 20 seconds, when the primary link is available and stabilized.

Which configuration resolves the issue?

A. R1(config)#track 700 ip sla 700 R1 (config-track)#delay down 30 up 20

B. R1 (config)#ip sla 700 R1(config-ip-sla)#delay down 30 up 20

C. R1 (config)#ip sla 700 R1(config-ip-sla)#delay down 20 up 30

D. R1(config)#track 700 ip sla 700 R1(config-track)#delay down 20 up 30

Correct Answer: *A*

Community vote distribution

A (100%)

forccnp 9 months, 1 week ago

Selected Answer: A

A is correct answer

upvoted 1 times

Noproblem22 1 year ago

A is correct

upvoted 1 times

WAKIDI 1 year, 5 months ago

Selected Answer: A

https://www.cisco.com/c/en/us/support/docs/smb/switches/cisco-550x-series-stackable-managed-switches/smb5797-configure-ip-sla-tracking-for-ipv4-static-routes-on-an-sg550.html#:~:text=To%20configure%20a%20period%20of%20time%20in%20seconds%20to%20delay%20state%20changes%20of%20a%20tracking%20object%2C%20enter%20the%20following%3A

upvoted 2 times

Question #210 *Topic 1*

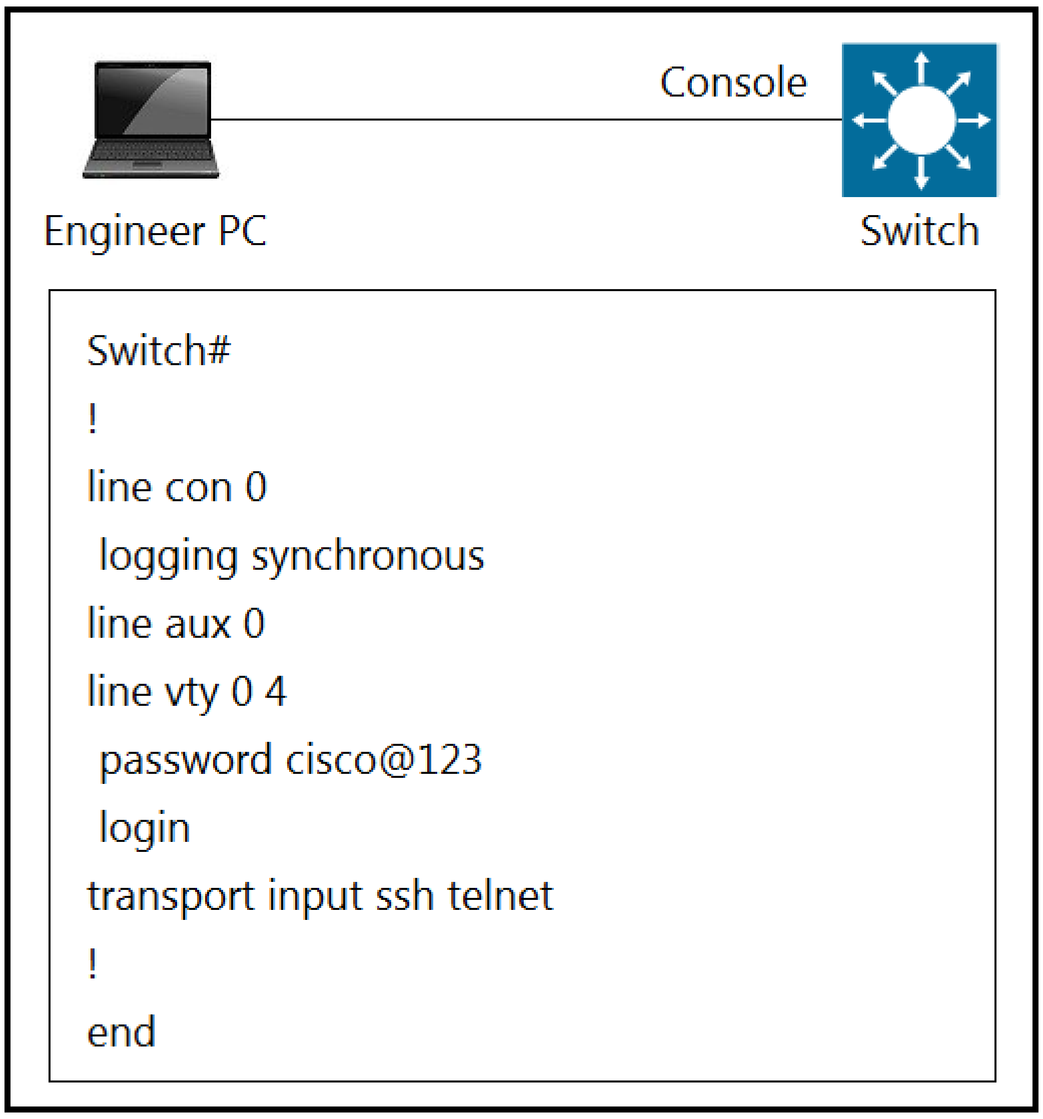

```
Switch#
!
line con 0
 logging synchronous
line aux 0
line vty 0 4
 password cisco@123
 login
transport input ssh telnet
!
end
```

Refer to the exhibit. An engineer must block access to the console ports for all corporate remote Cisco devices based on the recent corporate security policy but the security team still can connect through the console port.
Which configuration on the console port resolves the issue?

A. login and password

B. exec 0 0

C. transport input telnet

D. no exec

Correct Answer: *D*

Community vote distribution

D (100%)

Pietjeplukgeluk 1 week, 5 days ago
Selected answer D is correct, but please note "transport input none" would be a better solution in real life.
upvoted 1 times

inteldarvid 5 months, 2 weeks ago
Selected Answer: D
option correct is "D"
https://www.tenable.com/audits/items/CIS_Cisco_IOS_15_v4.0.1_Level_1.audit:f6d68c36cfcc77325b421f9865134f41
upvoted 1 times

IceFireSoul 1 year, 2 months ago
Provided answer is correct
For reference see:
https://community.cisco.com/t5/routing/no-exec/td-p/3715737
upvoted 2 times

Question #211 *Topic 1*

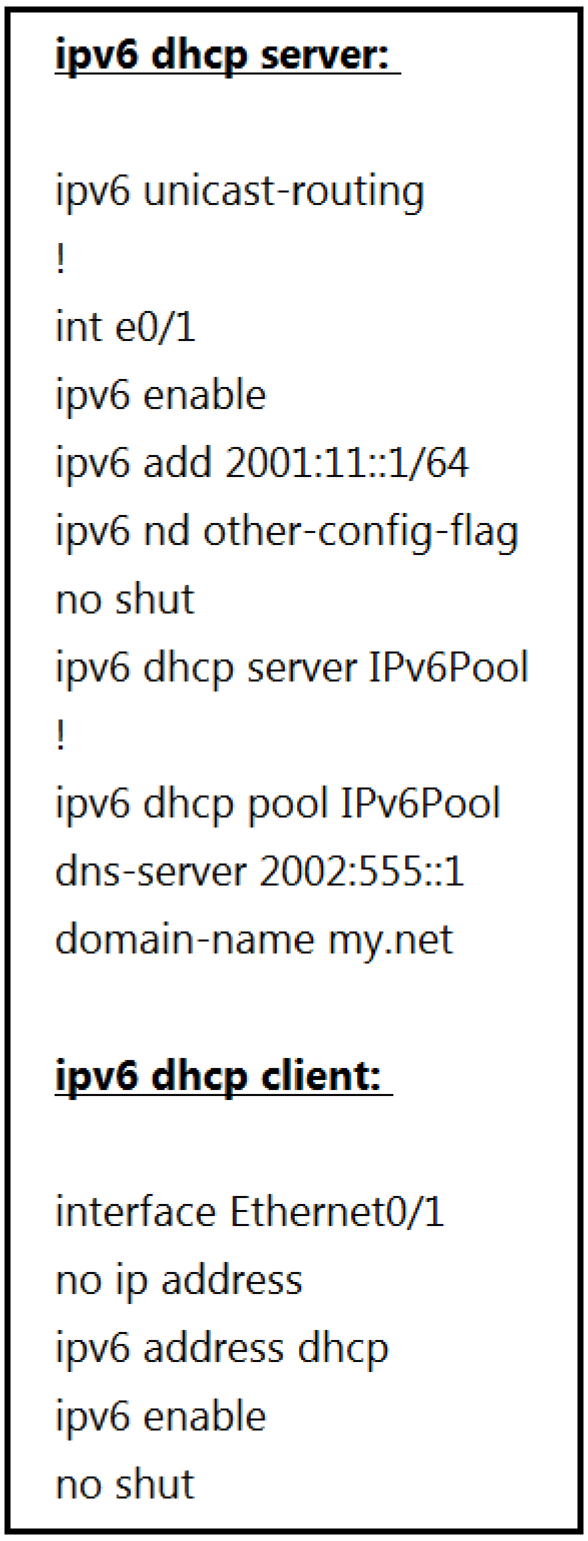

ipv6 dhcp server:

```
ipv6 unicast-routing
!
int e0/1
ipv6 enable
ipv6 add 2001:11::1/64
ipv6 nd other-config-flag
no shut
ipv6 dhcp server IPv6Pool
!
ipv6 dhcp pool IPv6Pool
dns-server 2002:555::1
domain-name my.net
```

ipv6 dhcp client:

```
interface Ethernet0/1
no ip address
ipv6 address dhcp
ipv6 enable
no shut
```

Refer to the exhibit. A network administrator is troubleshooting IPv6 address assignment for a DHCP client that is not getting an IPv6 address from the server.
Which configuration retrieves the client IPv6 address from the DHCP server?

A. ipv6 address autoconfig command on the interface

B. ipv6 dhcp server automatic command on DHCP server

C. ipv6 dhcp relay-agent command on the interface

D. service dhcp command on DHCP server

Correct Answer: *A*

Community vote distribution

A (100%)

⊟ **inteldarvid** 5 months, 2 weeks ago

Selected Answer: A

option A:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_dhcp/configuration/xe-16/dhcp-xe-16-book/ip6-dhcp-stateless-auto.html

upvoted 1 times

⊟ **inteldarvid** 5 months, 2 weeks ago

Selected Answer: A

option A is correct

upvoted 1 times

HungarianDish 7 months, 3 weeks ago

Selected Answer: A

The closest solution seems to be "ipv6 address autoconfig", however this configuration is not going to achieve ipv6 address assignment by the DHCPv6 server.

https://www.routerfreak.com/the-idiosyncrasies-of-ipv6-on-cisco-devices/
https://networklessons.com/ipv6/cisco-dhcpv6-server-configuration

upvoted 2 times

HungarianDish 7 months, 3 weeks ago

Based on the output, the ipv6 dhcp server is configured to do DHCPv6 Stateless Configuration.
The command "ipv6 nd other-config-flag" on the server tells the client to use DHCPv6 to receive extra information
(domain name and DNS server) after they used autoconfiguration.
However, the client is configured to use "ipv6 address dhcp" to obtain an address through stateful DHCPv6, that process is not working.
Stateful DHCPv6 (obtaining ipv6 address) is not possible since the ipv6 address prefix is not set under the dhcpv6 pool.

upvoted 2 times

HungarianDish 7 months, 3 weeks ago

After all, the client needs "ipv6 address autoconfig" to obtain an ipv6 address. -> Besides the link-local address, a global unicast address is also going to be been added (if there are other ipv6 devices on the segment).
"ipv6 enable" -> IPv6 is enabled on the interface, so the interface has been automatically configured with a link-local IPv6 address.
This command is not needed if we use "ipv6 address autoconfig" .

upvoted 1 times

MasterMatt 8 months ago

This question is unclear. "Which configuration retrieves the client IPv6 address from the DHCP server?" This automatically raises questions as to which part we are required to configure. From the client output we only have port level configuration. So from all the valid options the "ipv6 address autoconfig" is the one that match to what we have. However with this command you configure SLAAC and not DHCP.

upvoted 2 times

Nhan 1 year, 6 months ago

A is correct, the first step is using the ipv6 address autoconfig to create the link-local address for the interface, then the ip address will be assign to the interface using dhcp server, and the command is "ipv6 address dhcp"

upvoted 1 times

piojo 1 year, 6 months ago

Selected Answer: A

The config is invalid, there is no such "ipv6 address dhcp" command for the client.

upvoted 1 times

JingleJangus 1 year, 5 months ago

Seems you might be wrong on this one:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/command/ipv6-cr-book/ipv6-i1.html#wp2212047392

upvoted 2 times

Question #212 Topic 1

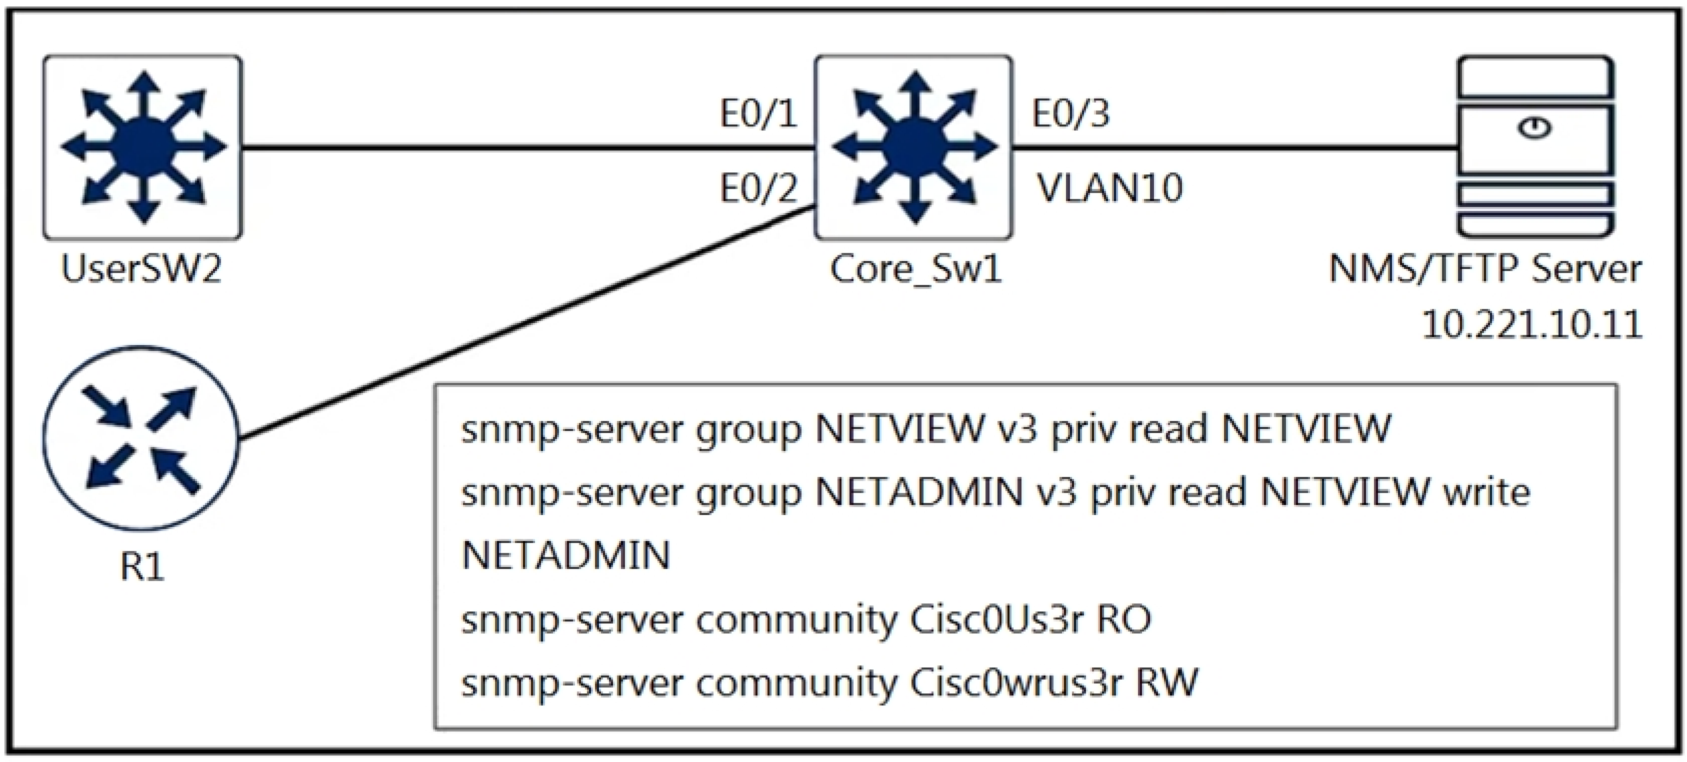

```
snmp-server group NETVIEW v3 priv read NETVIEW
snmp-server group NETADMIN v3 priv read NETVIEW write
NETADMIN
snmp-server community Cisc0Us3r RO
snmp-server community Cisc0wrus3r RW
```

Refer to the exhibit. A junior engineer configured SNMP to network devices. Malicious users have uploaded different configurations to the network devices using
SNMP and TFTP servers.
Which configuration prevents changes from unauthorized NMS and TFTP servers?

A. access-list 20 permit 10.221.10.11 access-list 20 deny any log ! snmp-server group NETVIEW v3 priv read NETVIEW access 20 snmp-server group NETADMIN v3 priv read NETVIEW write NETADMIN access 20 snmp-server community Cisc0Us3r RO 20 snmp-server community Cisc0wrus3r RW 20 snmp-server tftp-server-list 20

B. access-list 20 permit 10.221.10.11 access-list 20 deny any log ! snmp-server group NETVIEW v3 priv read NETVIEW access 20 snmp-server group NETADMIN v3 priv read NETVIEW write NETADMIN access 20 snmp-server community Cisc0wrus3r RO 20 snmp-server community Cisc0Us3r RW 20 snmp-server tftp-server-list 20

C. access-list 20 permit 10.221.10.11 access-list 20 deny any log

D. access-list 20 permit 10.221.10.11

Correct Answer: *A*

Community vote distribution

A (100%)

mrnipsnips Highly Voted 1 year, 1 month ago

Man cisco are petty AF

upvoted 5 times

Slinky 10 months ago

Absolutely died laughing at this but it's true

upvoted 2 times

ledesir 2 weeks, 3 days ago

hahhahhhha same thing for me

upvoted 1 times

ZamanR Most Recent 5 days, 20 hours ago

A is correct answer

upvoted 1 times

Jey117 2 months, 1 week ago

Are you kidding? You can fail this question just because they inverted communities? Cisco WTHell. Stop trying to take people's money. LOL

upvoted 1 times

Colmenarez 3 months, 3 weeks ago

Spot the difference type of question hahahaha

upvoted 2 times

MasterMatt 8 months ago

Selected Answer: A

access-list 20 permit 10.221.10.11 --> Permitting only from NMS.
access-list 20 deny any log --> Similar to implicit deny by logging is enabled.

snmp-server group NETVIEW v3 priv read NETVIEW access 20 --> We filter based on the access-list
snmp-server group NETADMIN v3 priv read NETVIEW write NETADMIN access 20 --> We filter based on the access-list

snmp-server community Cisc0Us3r RO 20 --> Same level of permission but we filter based on the access-list
snmp-server community Cisc0wrus3r RW 20 --> Same level of permission but we filter based on the access-list
snmp-server tftp-server-list 20 --> Limit TFTP servers used via SNMP only over access-list 20

upvoted 3 times

JOKERR 1 year, 7 months ago

Isn't t the answer B?

Because B has the RW community string...

upvoted 2 times

Bolt_Action_Studios 1 year, 7 months ago

Community strings are reversed with B

upvoted 2 times

Question #213 *Topic 1*

An engineer creates a Cisco DNA Center cluster with three nodes, but all the services are running on one host node.
Which action resolves this issue?

A. Restore the link on the switch interface that is connected to a cluster link on the Cisco DNA Center.

B. Click system updates, and upgrade to the latest version of Cisco DNA Center.

C. Enable service distribution from the Systems 360 page.

D. Click the master host node with all the services and select services to be moved to other hosts.

Correct Answer: *C*

Community vote distribution

C (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: C

C correct:

https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/1-3-3-0/ha_guide/b_cisco_dna_center_ha_guide_1_3_3_0.html

upvoted 1 times

Pbshah 1 year, 5 months ago

Selected Answer: C

Answer C is correct

upvoted 1 times

xziomal9 1 year, 7 months ago

Selected Answer: C

The correct answer is: C

upvoted 1 times

Bruffas 1 year, 7 months ago

Selected Answer: C

C.
Click and then choose System Settings.
The System 360 tab is displayed by default.
2. In the Hosts area, click Enable Service Distributio

upvoted 2 times

Question #214

Topic 1

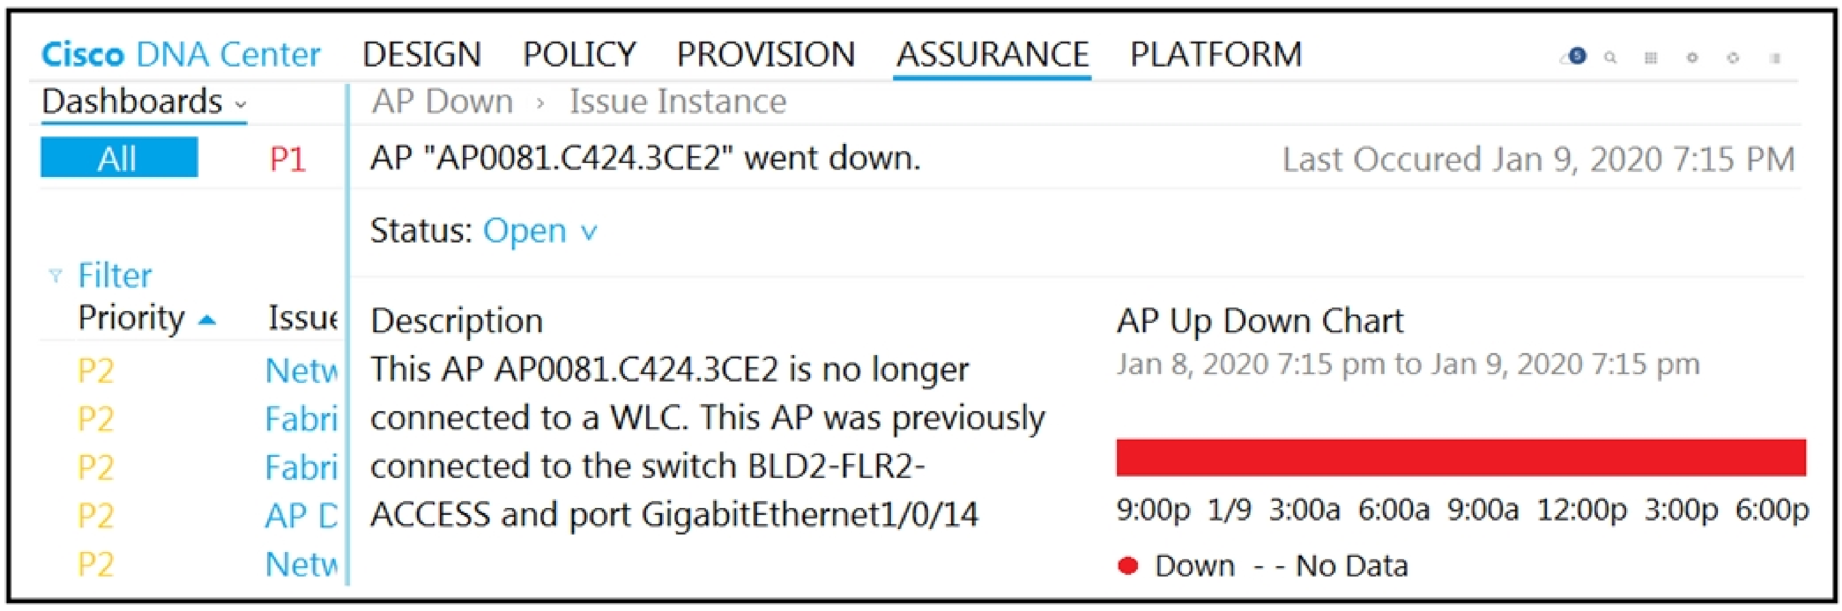

AP Down › Issue Instance

AP "AP0081.C424.3CE2" went down.

Last Occured Jan 9, 2020 7:15 PM

Status: Open v

Description

This AP AP0081.C424.3CE2 is no longer connected to a WLC. This AP was previously connected to the switch BLD2-FLR2-ACCESS and port GigabitEthernet1/0/14

AP Up Down Chart

Jan 8, 2020 7:15 pm to Jan 9, 2020 7:15 pm

Refer to the exhibit. The AP status from Cisco DNA Center Assurance Dashboard shows some physical connectivity issues from access switch interface G1/0/14.

Which command generates the diagnostic data to resolve the physical connectivity issues?

A. check cable-diagnostics tdr interface GigabitEthernet1/0/14

B. verify cable-diagnostics tdr interface GigabitEthernet1/0/14

C. show cable-diagnostics tdr interface GigabitEthernet1/0/14

D. test cable-diagnostics tdr interface GigabitEthernet1/0/14

Correct Answer: *C*

Community vote distribution

D (100%)

Mishranihal737 2 months, 3 weeks ago

Selected Answer: D

Yes , first u need to run that test command to generate data and then use show command to view.

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

option D is corerct, i test in my CORE :)

upvoted 1 times

PimplePooper 12 months ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

NoUserName1234 1 year, 2 months ago

D is correct. The question is how to generate the output, the only way to do this is by the global command test cable diagnostic.

upvoted 3 times

maewzilla 1 year, 4 months ago

D. test cable-diagnostics tdr generates
result.

upvoted 2 times

TECH3K3 1 year, 5 months ago

Selected Answer: D

The answer is D, as I use this command often when at work

upvoted 2 times

Orchidium 1 year, 5 months ago

I would personally go with D "test cable-diagnostics tdr interface GigabitEthernet1/0/14" since the question asks what command will "generate" (not "display the output of") the diagnostic data needed

upvoted 3 times

Nhan 1 year, 6 months ago

the question is "Which command generates the diagnostic data", I think the answer D is more relevant

upvoted 2 times

xziomal9 1 year, 7 months ago

Selected Answer: D

The correct answer is: D

upvoted 2 times

jthompaf 1 year, 7 months ago

I feel like this is poorly written. Test cable-diagnostics, generates the information, but show cable-diagnostics actually show the output of the test. Can someone clarify what the answer should be?

upvoted 4 times

Question #215 *Topic 1*

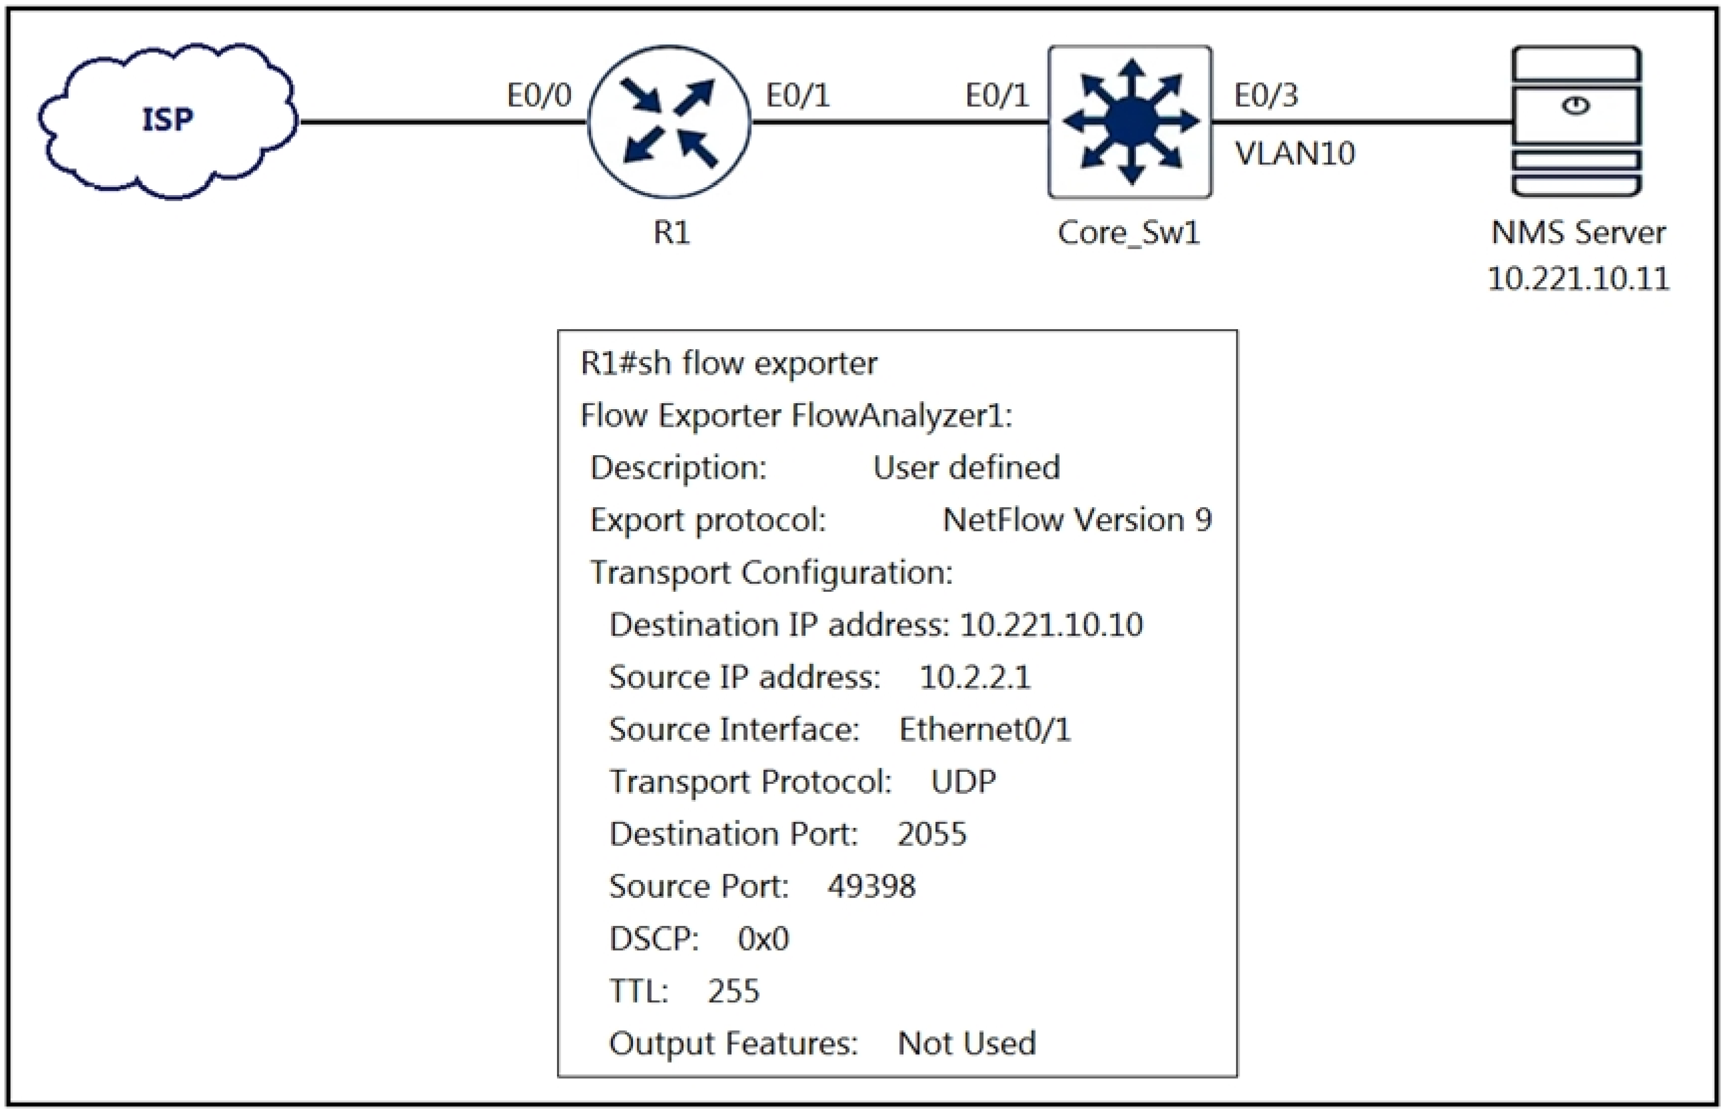

```
R1#sh flow exporter
Flow Exporter FlowAnalyzer1:
 Description:        User defined
 Export protocol:         NetFlow Version 9
 Transport Configuration:
   Destination IP address: 10.221.10.10
   Source IP address:    10.2.2.1
   Source Interface:   Ethernet0/1
   Transport Protocol:   UDP
   Destination Port:   2055
   Source Port:   49398
   DSCP:   0x0
   TTL:   255
   Output Features:   Not Used
```

Refer to the exhibit. An engineer configured NetFlow on R1, but the NMS server cannot see the flow from R1.
Which configuration resolves the issue?

A. interface Ethernet0/1 flow-destination 10.221.10.11

B. interface Ethernet0/0 flow-destination 10.221.10.11

C. flow exporter FlowAnalyzer1 destination 10.221.10.11

D. flow monitor Flowmonitor1 destination 10.221.10.11

Correct Answer: *C*

Community vote distribution

C (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: C

option C is correct

https://www.cisco.com/c/dam/en/us/td/docs/security/stealthwatch/netflow/Cisco_NetFlow_Configuration.pdf

upvoted 1 times

encor01 10 months ago

The given answer seems correct.

https://www.cisco.com/c/en/us/td/docs/iosxml/ios/fnetflow/configuration/15-mt/fnf-15-mt-book/cfg-de-fnflow-exprts.html

upvoted 1 times

chris7890 1 year, 2 months ago

Why not answer D? Since this is the primary connection?

upvoted 1 times

WAKIDI 1 year, 5 months ago

C seems ok. for A and B : there is no such command

upvoted 1 times

Question #216 Topic 1

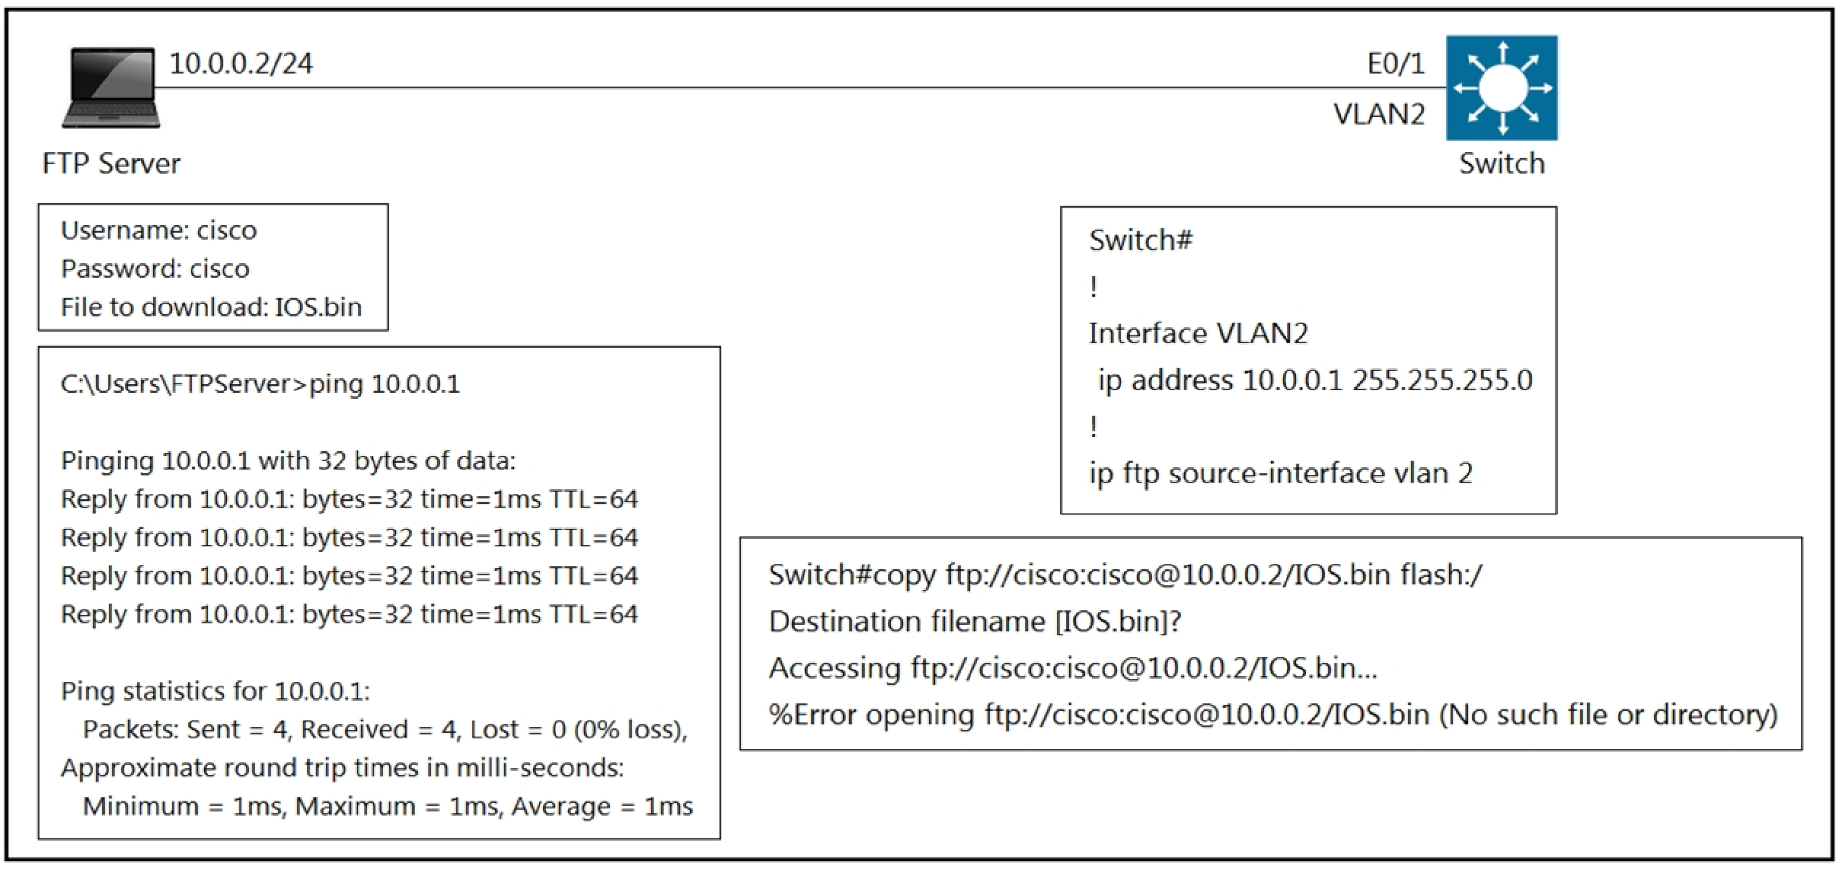

```
Username: cisco
Password: cisco
File to download: IOS.bin
```

```
C:\Users\FTPServer>ping 10.0.0.1

Pinging 10.0.0.1 with 32 bytes of data:
Reply from 10.0.0.1: bytes=32 time=1ms TTL=64
Reply from 10.0.0.1: bytes=32 time=1ms TTL=64
Reply from 10.0.0.1: bytes=32 time=1ms TTL=64
Reply from 10.0.0.1: bytes=32 time=1ms TTL=64

Ping statistics for 10.0.0.1:
   Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
   Minimum = 1ms, Maximum = 1ms, Average = 1ms
```

```
Switch#
!
Interface VLAN2
 ip address 10.0.0.1 255.255.255.0
!
ip ftp source-interface vlan 2
```

```
Switch#copy ftp://cisco:cisco@10.0.0.2/IOS.bin flash:/
Destination filename [IOS.bin]?
Accessing ftp://cisco:cisco@10.0.0.2/IOS.bin...
%Error opening ftp://cisco:cisco@10.0.0.2/IOS.bin (No such file or directory)
```

Refer to the exhibit. An engineer cannot copy the IOS.bin file from the FTP server to the switch.
Which action resolves the issue?

A. Allow file permissions to download the file from the FTP server.

B. Add the IOS.bin file, which does not exist on FTP server.

C. Make memory space on the switch flash or USB drive to download the file.

D. Use the copy flash:/ ftp://cisco@10.0.0.2/IOS.bin command.

Correct Answer: *B*

Community vote distribution

B (71%) A (29%)

inteldarvid 5 months, 1 week ago

Selected Answer: B

B si super correct:

Team look this:
https://quickview.cloudapps.cisco.com/quickview/bug/CSCeh27229

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

the option correct is B, bceause in the script, the admin put user and password (cisco:cisco). Its not is necesary put password and user in the switch (ip ftp user, ip ftp password)

upvoted 1 times

Malasxd 7 months, 2 weeks ago

When the user doesn't have permission to access a directory or file and when the file doesn't exist the error shown are the same (No such file or directory).

If the file existe, and I believe it existe because the file is shown bellow the username and password e answer "A" is corret. If the file doesn't exist the corret is "B".

I would chose "A", but it can be wrong.

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

This is a long thread, but it points out that missing permission is indicated by the error message: "Permission denied"

"No such file or directory" means the file and/or directory is not found in the specified directory of the TFTP server.
"Permission denied" means read access to the file and/or directory is not enabled.
https://community.cisco.com/t5/switching/error-opening-tftp-permission-denied/td-p/3302909/page/3

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Or the IOS.bin file might be under a different directory on the ftp server, and then still answer "B" is OK.

upvoted 2 times

HungarianDish 7 months, 3 weeks ago

Selected Answer: B

https://bst.cisco.com/bugsearch/bug/CSCeh27229

https://community.cisco.com/t5/switching/copy-flash-tftp-command-failed-on-cisco-3750-switch/td-p/1526415

The file does not exist under the specified directory on the ftp server. Solution "B".

upvoted 2 times

SujanSikrikar 10 months ago

https://community.cisco.com/t5/routing/can-t-copy-a-file-form-ftp-to-flash/td-p/821267
Correct answer is A.

```
switch# config t
switch(config)# ip ftp username cisco
switch(config)# ip ftp password cisco123
switch#copy ftp://cisco:cisco123@ftpserver//iosdirectory/ios_filename.bin slot0:ios_filename.bin
```

upvoted 1 times

HungarianDish 7 months, 3 weeks ago

The mentioned cisco article concludes that the error is rather due to an incorrect file name or location.

upvoted 1 times

Lilienen 10 months ago

Selected Answer: B

Correct answer: Add the IOS.bin file, which does not exist on FTP server.

upvoted 1 times

tseen 10 months, 2 weeks ago

Selected Answer: A

The error from the switch shows that it cannot open or find the file from the FTP server, hence the FTP server needs to grant permissions

upvoted 2 times

Nhan 1 year, 6 months ago

the error is clearly indicating that is no such file or directory, the given answer is correct

upvoted 2 times

Question #217 *Topic 1*

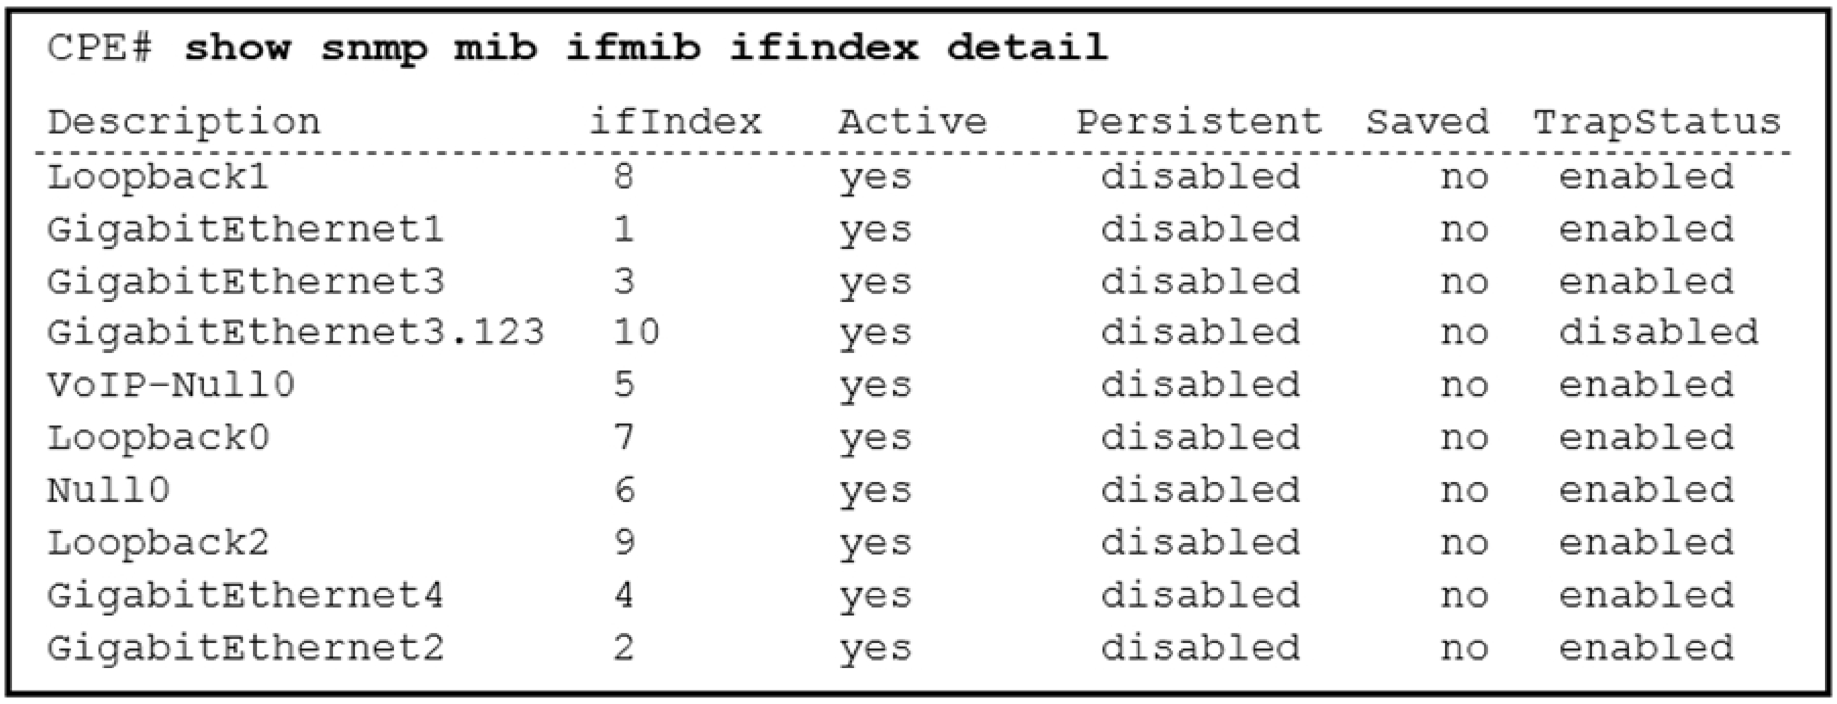

```
CPE# show snmp mib ifmib ifindex detail

Description            ifIndex   Active   Persistent  Saved  TrapStatus
-----------------------------------------------------------------------
Loopback1                8        yes      disabled     no    enabled
GigabitEthernet1         1        yes      disabled     no    enabled
GigabitEthernet3         3        yes      disabled     no    enabled
GigabitEthernet3.123     10       yes      disabled     no    disabled
VoIP-Null0               5        yes      disabled     no    enabled
Loopback0                7        yes      disabled     no    enabled
Null0                    6        yes      disabled     no    enabled
Loopback2                9        yes      disabled     no    enabled
GigabitEthernet4         4        yes      disabled     no    enabled
GigabitEthernet2         2        yes      disabled     no    enabled
```

Refer to the exhibit. After reloading the router, an administrator discovered that the interface utilization graphs displayed inconsistencies with their previous history in the NMS.

Which action prevents this issue from occurring after another router reload in the future?

A. Configure SNMP interface index persistence on the router.

B. Save the router configuration to startup-config before reloading the router.

C. Rediscover all the router interfaces through SNMP after the router is reloaded.

D. Configure SNMP to use static OIDs referring to individual router interfaces.

Correct Answer: *A*

Reference:

https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/15S/configuration/guide/7600_15_0s_book/ifindx.pdf

Community vote distribution

A (100%)

HungarianDish 7 months, 3 weeks ago

Selected Answer: A

Answer and provided source are correct.
https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/15S/configuration/guide/7600_15_0s_book/ifindx.pdf
https://packetlife.net/blog/2010/apr/22/snmp-interface-index-persistence/

upvoted 3 times

Question #218 Topic 1

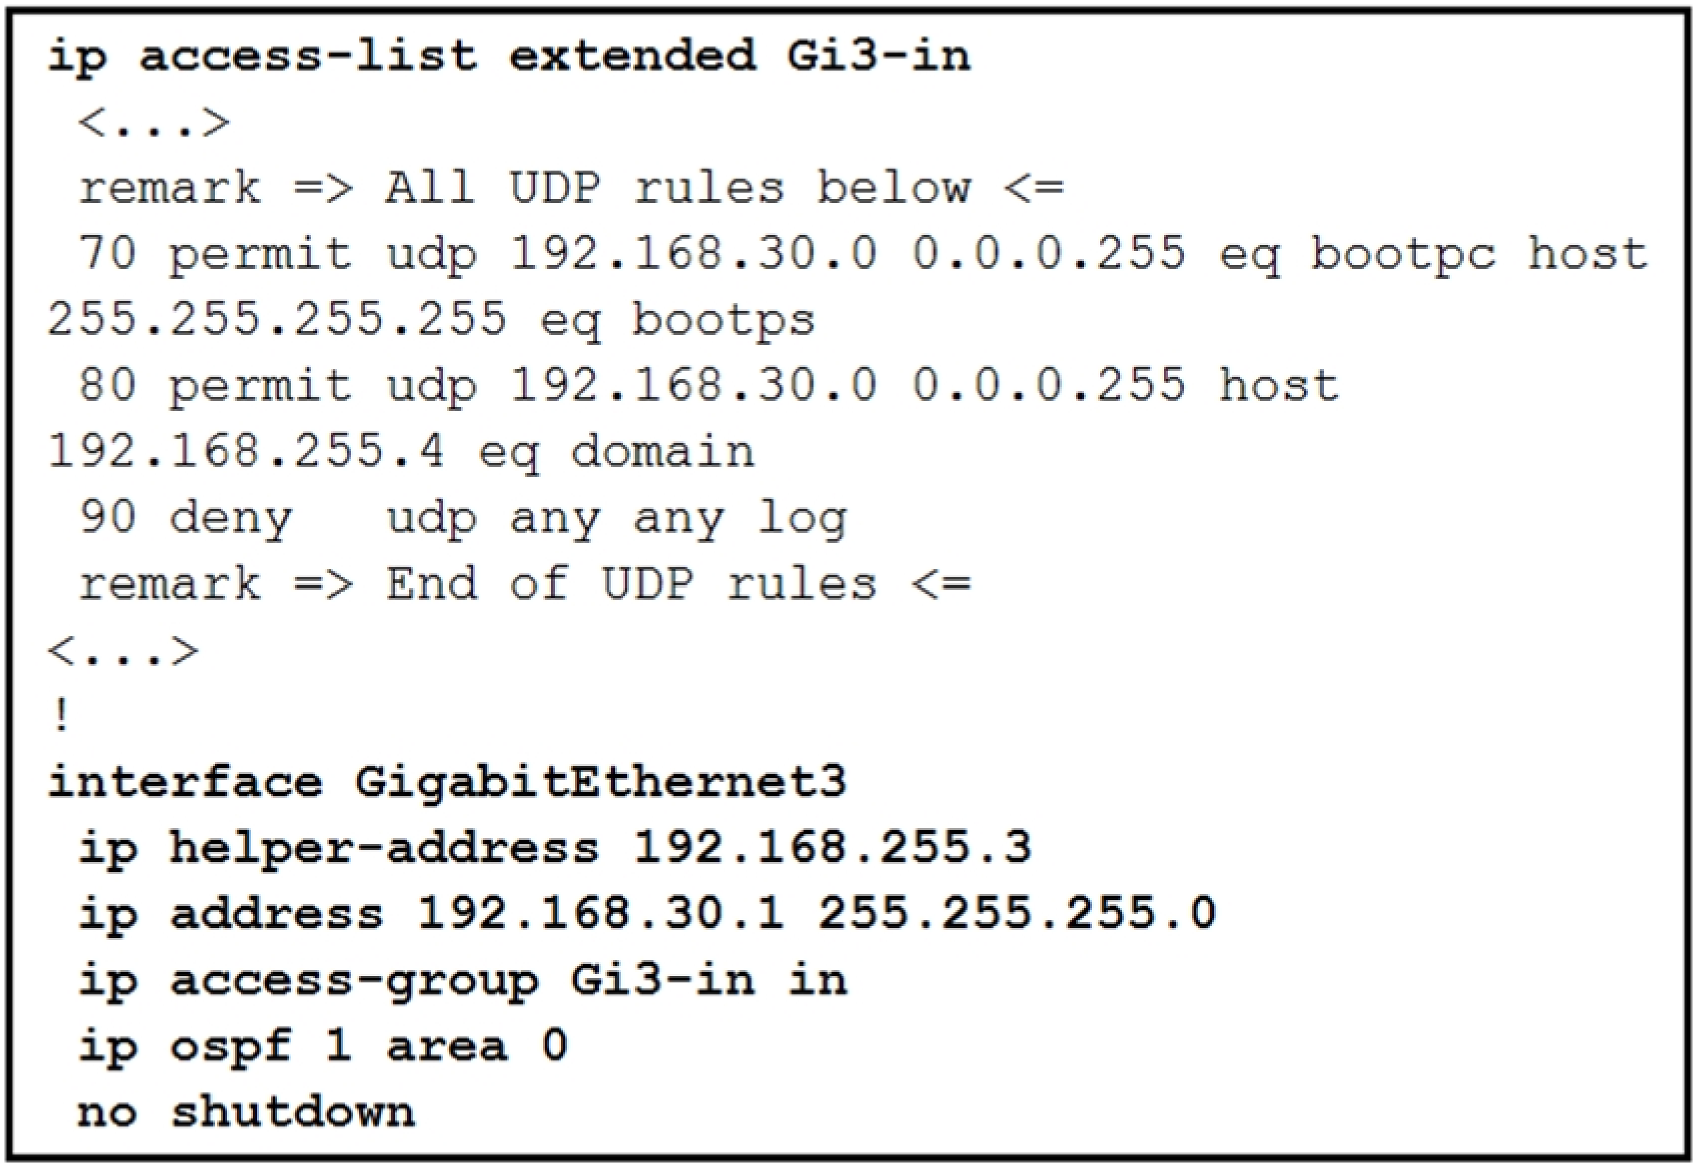

```
ip access-list extended Gi3-in
 <...>
 remark => All UDP rules below <=
 70 permit udp 192.168.30.0 0.0.0.255 eq bootpc host
255.255.255.255 eq bootps
 80 permit udp 192.168.30.0 0.0.0.255 host
192.168.255.4 eq domain
 90 deny   udp any any log
 remark => End of UDP rules <=
<...>
!
interface GigabitEthernet3
 ip helper-address 192.168.255.3
 ip address 192.168.30.1 255.255.255.0
 ip access-group Gi3-in in
 ip ospf 1 area 0
 no shutdown
```

Refer to the exhibit. In an attempt to increase the network security, the administrator applied the Gi3-in ACL to the Gi3 interface. After the ACL was applied, clients in the network connected to Gi3 lost their ability to obtain IP settings from DHCP.
Which two configuration commands must be added to the Gi3-in ACL to reinstate the DHCP service for the clients? (Choose two.)

A. 74 permit udp 192.168.30.0 0.0.0.255 eq bootpc host 192.168.255.3 eq bootps

B. 71 permit udp host 0.0.0.0 eq bootps host 255.255.255.255 eq bootpc

C. 73 permit udp host 0.0.0.0 eq bootpc host 192.168.255.3 eq bootps

D. 72 permit udp host 192.168.255.3 eq bootps 192.168.30.0 0.0.0.255 eq bootpc

E. 75 permit udp host 0.0.0.0 eq bootpc host 255.255.255.255 eq bootps

Correct Answer: *B*
Reference:
https://community.spiceworks.com/topic/1982739-help-with-access-list-to-permit-dhcp-requests-and-renews

Community vote distribution
A (60%) B (20%) 13% 7%

Huntkey Highly Voted 1 year, 2 months ago
For first time DHCP client, the discover and request messages would all be from 0.0.0.0 to 255.255.255.255. So E is needed.
https://wiki.wireshark.org/uploads/__moin_import__/attachments/DHCP/dhcp-ws.png
For renewing with DHCP request, the source is the current assigned IP and the destination is server itself. So A is needed.
Other packets like inform is from the assigned IP to the 255.255.255.255. The existing ACL entry allows for it already.
I will go with AE.
upvoted 11 times

guy276465281819372 Most Recent 4 months, 2 weeks ago
Selected Answer: E
A & E CORRECT
upvoted 1 times

inteldarvid 5 months, 2 weeks ago
Selected Answer: A
A and E correct

https://networkengineering.stackexchange.com/questions/38044/dhcp-bootpc-acl

upvoted 2 times

Malasxd 7 months, 1 week ago

Selected Answer: A

A and E are correct.

upvoted 2 times

HungarianDish 7 months, 3 weeks ago

Selected Answer: E

For me E + A.
https://community.cisco.com/t5/switching/acl-not-working-as-intended/td-p/4168422
"permit udp host 0.0.0.0 eq boopc host 255.255.255.255 eq bootps. For the DHCP IP renewal, you can configure permit udp 10.20.20.0 0.0.0.255 eq bootpc host 128.1.99.1 eq bootps. Reason why the one you configured would not work for DHCP DORA is because when the client first time tries to get an IP, it sources with 0.0.0.0, and the DHCP request will be broadcasted to the IP 255.255.255.255. However, when the client tries to renew its IP address, it would source from its IP address which will be within the subnet 10.20.20.0/24, and will send the renewal request to the DHCP server IP as unicast."
https://www.certforums.com/threads/acl-allow-access-to-dhcp-server.36762/

upvoted 1 times

6dd4aa0 8 months, 2 weeks ago

Selected Answer: D

From Figure 7-2 (Pg 127 CCNA 200-301 Volume 2)
DHCP Client
PC-A --(From: 0.0.0.0 To: 255.255.255.255)--> Router ----------------------------------> DHCP Server
192.168.30.1 192.168.255.3
Option E will be correct.
===
From Figure 7-3 (Pg 128 CCNA 200-301 Volume 2)
PC-A <------------------------- Router <---(From 192.168.255.3 To:192.168.30.1)--- DHCP Server
192.168.30.1 192.168.255.3
Option D will be correct.

upvoted 1 times

pyrokar 7 months, 1 week ago

It is an inbound ACL, it does not filter answers from the server

upvoted 1 times

pyrokar 7 months, 1 week ago

To be more precise, it is inbound on the interface facing the clients. Since there is a helper-address configured (in another subnet), the dhcp server is on another interface. So this ACL is not applied to answeers from the server.
It would if it was applied outbound or on another interface.

upvoted 1 times

Typovy 9 months, 1 week ago

A E is correct.
D is pointless because source IP address is DHCP server addres. This ACL is applied to LAN facing interface inbound so DHCP server as source here will have no matches :)

upvoted 2 times

TECH3K3 1 year, 4 months ago

I labbed this and none of the combinations worked.
B, C and D can't be added to the ACL as I get a message "% % Duplicate sequence number."
A and E can be added to the ACL but the PC doesn't get an IP address and you get a syslog message ... list Gi3-in denied udp 192.168.30.2(67) -> 192.168.30.100(68), 2 packets
If I removed the ACL the PC gets an IP address

upvoted 1 times

TECH3K3 1 year, 4 months ago

UPDATE!!
I moved the ACL from a router interface to a swicth interface and ONLY E was needed for me to obtain an IP address

upvoted 1 times

TECH3K3 1 year, 4 months ago

I also think it's D and E just from looking at past configs for a company I use to work for

upvoted 2 times

johnu329 1 year, 5 months ago

D and E

bootpc = udp/68
bootps = udp/67

Client-end port is 68; Server-end port is 67 (http://klamp.works/2016/04/29/dhcp.html)

Therefore, correct answers are D and E:
--> To server (port 67)
E. 75 permit udp host 0.0.0.0 eq bootpc host 255.255.255.255 eq bootps

--> To client (port 68)
D. 72 permit udp host 192.168.255.3 eq bootps 192.168.30.0 0.0.0.255 eq bootpc
upvoted 2 times

JingleJangus 1 year, 6 months ago

Selected Answer: A

A and E
To get this question, you MUST be comfortable with the DHCP-DORA Exchange.
Discover:
Src: 0.0.0.0
Dest: 255.255.255.255

Offer:
Src: <DHCP Server Address>
Dest: <Relay Address> OR 255.255.255.255

Request:
Src: 0.0.0.0
Dest: 255.255.255.255

Ack:
Src: <DHCP Server Address>
Dest: <Relay Address> OR 255.255.255.255

Given the Inbound ACL applied to the Client-Facing Interface, AT A MINIMUM, E is required.
DHCP will also use Unicast for other operations and upkeep, so A is also important.
https://community.cisco.com/t5/switching/concerning-acl-with-dhcp/td-p/1239487
upvoted 4 times

t1s 1 year ago

Yes, A & E is correct.
E > for DORA
A > for renew
https://www.cloudshark.org/captures/0009d5398f37
upvoted 2 times

piojo 1 year, 6 months ago

Selected Answer: A

LABED it. Correct are A and B.

It should be FROM bootpc (client) TO bootps (server).

Source is 0.0.0.0 to 255.255.255.255 when first get and IP
Source is 192.168.30.X to 192.168.30.3 when renewing.
upvoted 1 times

WAKIDI 1 year, 5 months ago

did you mean A and E ?. the usage of bootc and boots seems to be better in E.
upvoted 2 times

piojo 1 year, 6 months ago

Sorry, A and C.
upvoted 1 times

JingleJangus 1 year, 6 months ago

I would disagree;
Clients initially send to a Broadcast Destination of 255.255.255.255, not Unicast.
Yes, the Relay is going to modify the Destination to Unicast; but since the ACL is applied in the inbound direction, this Destination translation is only going to happen AFTER the ACL has been applied to received traffic.
https://community.cisco.com/t5/switching/concerning-acl-with-dhcp/td-p/1239487
upvoted 2 times

xziomal9 1 year, 7 months ago

Selected Answer: B

The correct answer is: BE
B. 71 permit udp host 0.0.0.0 eq bootps host 255.255.255.255 eq bootpc
E. 75 permit udp host 0.0.0.0 eq bootpc host 255.255.255.255 eq bootps
upvoted 3 times

xziomal9 1 year, 7 months ago

DHCP uses 2 ports .UDP port number 67 is the destination port of a server, and UDP port number 68 is used by the client. (aka bootps and bootpc) DHCP process is abriviated as DORA.
upvoted 1 times

jthompaf 1 year, 7 months ago

All DHCP Discovers and Request received from the 192.168.30.0 network on the interface will be sent from 0.0.0.0 with a destination of 255.255.255.255, respectively. Therefore, only access-list with 0.0.0.0 and 255.255.255.255 (broadcast) addresses are relevant. E seems to be the only command necessary to get clients to pull DHCP information and to bind in the DHCP server.

upvoted 2 times

Ash78 1 year, 7 months ago

Hi, Please can anyone explain this?
Thank you

upvoted 2 times

Kimaf 1 year, 8 months ago

A & D are the right answers.

upvoted 2 times

Question #219 Topic 1

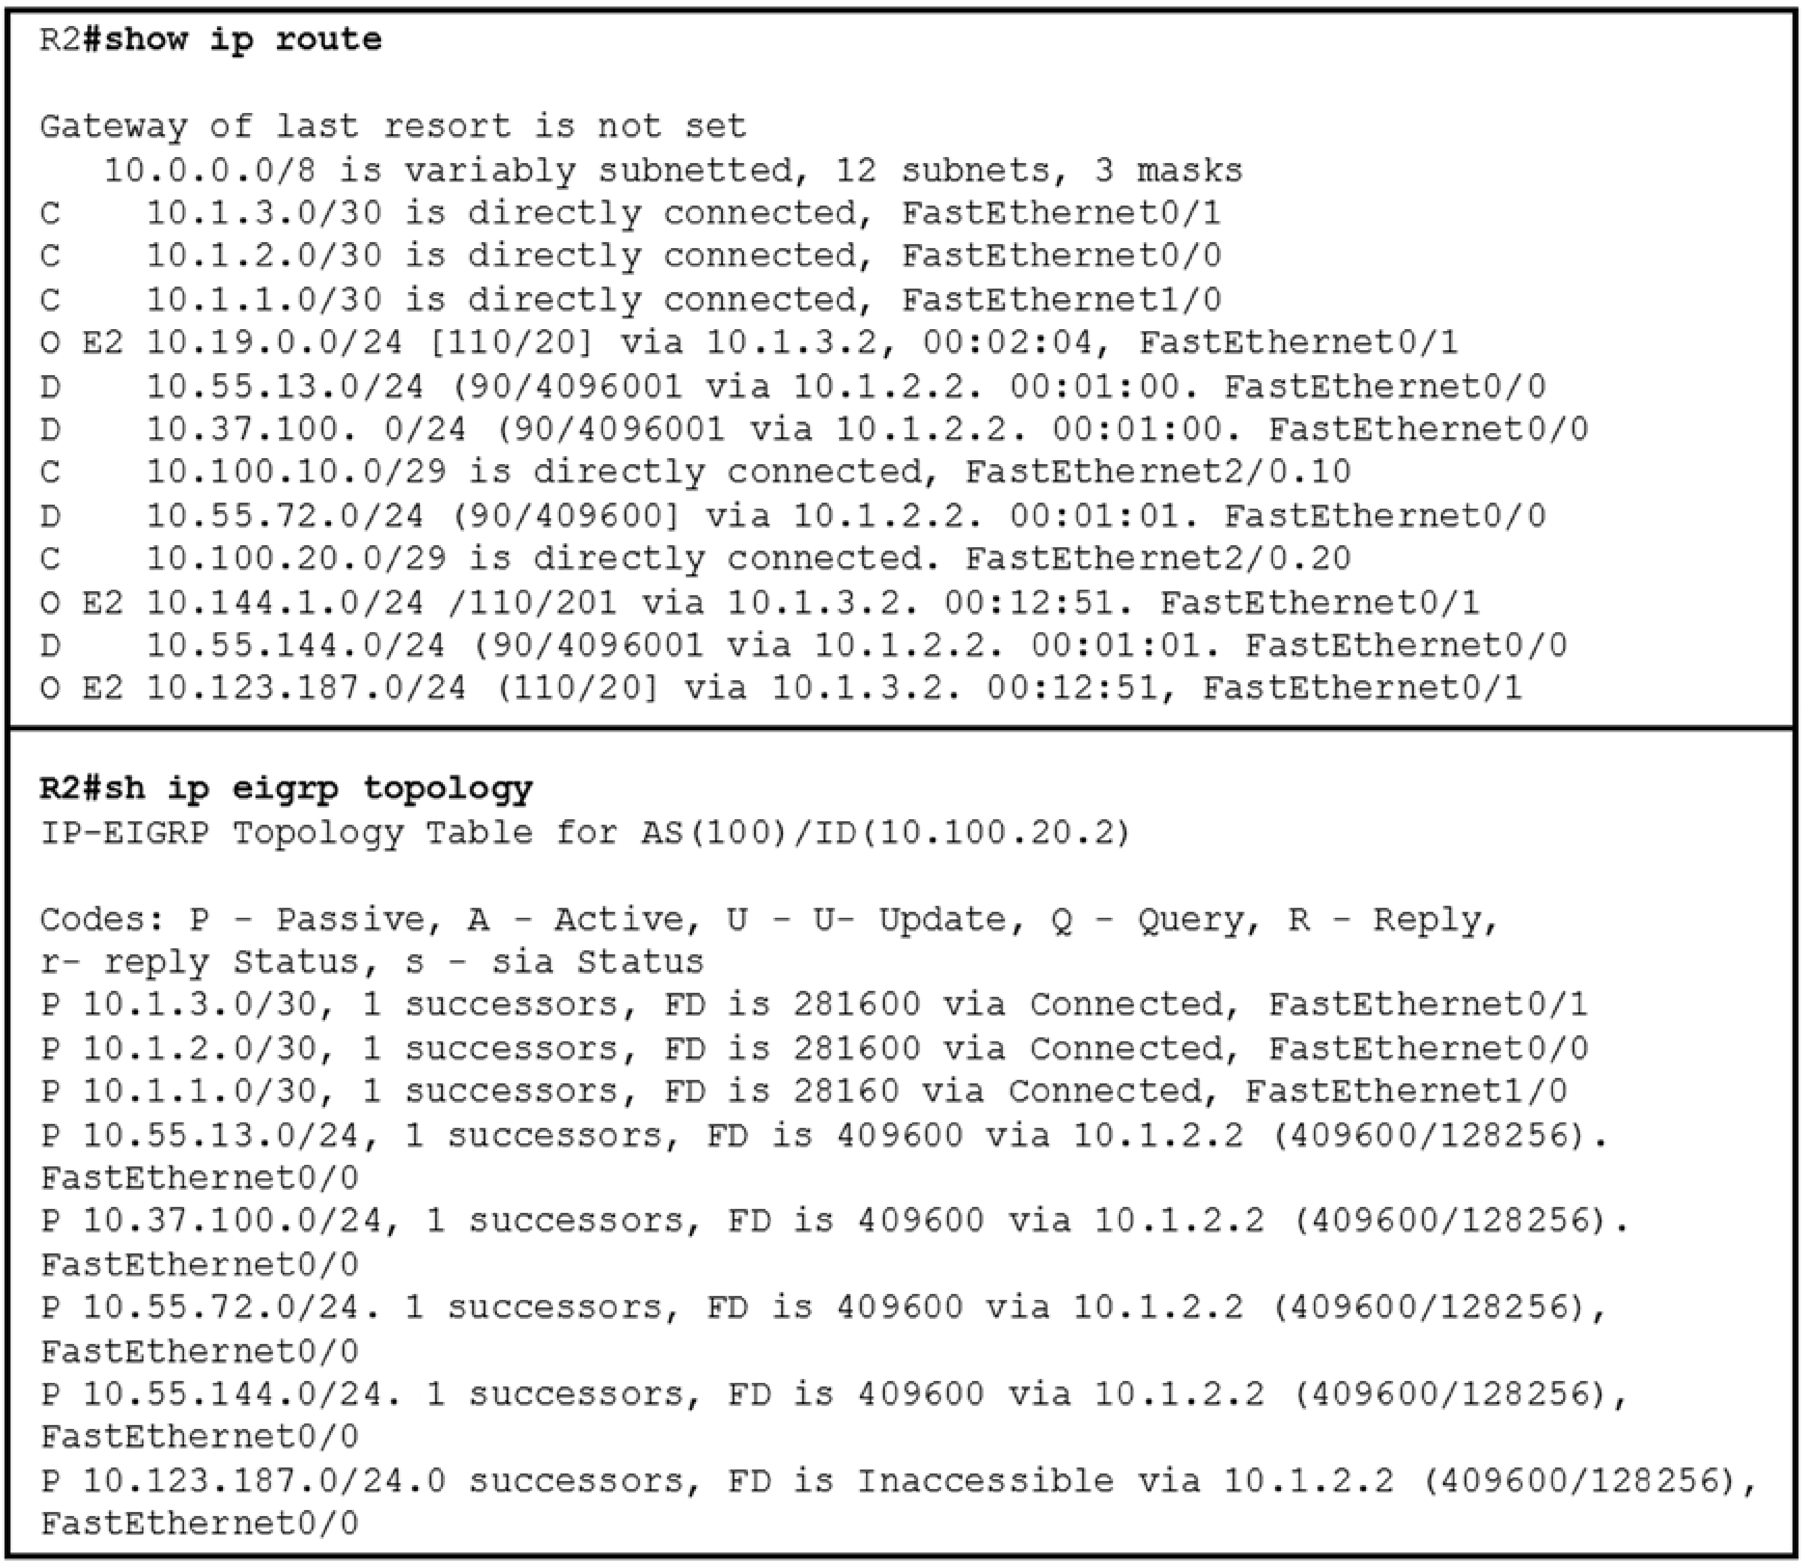

```
R2#show ip route

Gateway of last resort is not set
   10.0.0.0/8 is variably subnetted, 12 subnets, 3 masks
C    10.1.3.0/30 is directly connected, FastEthernet0/1
C    10.1.2.0/30 is directly connected, FastEthernet0/0
C    10.1.1.0/30 is directly connected, FastEthernet1/0
O E2 10.19.0.0/24 [110/20] via 10.1.3.2, 00:02:04, FastEthernet0/1
D    10.55.13.0/24 (90/4096001 via 10.1.2.2. 00:01:00. FastEthernet0/0
D    10.37.100. 0/24 (90/4096001 via 10.1.2.2. 00:01:00. FastEthernet0/0
C    10.100.10.0/29 is directly connected, FastEthernet2/0.10
D    10.55.72.0/24 (90/409600] via 10.1.2.2. 00:01:01. FastEthernet0/0
C    10.100.20.0/29 is directly connected. FastEthernet2/0.20
O E2 10.144.1.0/24 /110/201 via 10.1.3.2. 00:12:51. FastEthernet0/1
D    10.55.144.0/24 (90/4096001 via 10.1.2.2. 00:01:01. FastEthernet0/0
O E2 10.123.187.0/24 (110/20] via 10.1.3.2. 00:12:51, FastEthernet0/1
```

```
R2#sh ip eigrp topology
IP-EIGRP Topology Table for AS(100)/ID(10.100.20.2)

Codes: P - Passive, A - Active, U - U- Update, Q - Query, R - Reply,
r- reply Status, s - sia Status
P 10.1.3.0/30, 1 successors, FD is 281600 via Connected, FastEthernet0/1
P 10.1.2.0/30, 1 successors, FD is 281600 via Connected, FastEthernet0/0
P 10.1.1.0/30, 1 successors, FD is 28160 via Connected, FastEthernet1/0
P 10.55.13.0/24, 1 successors, FD is 409600 via 10.1.2.2 (409600/128256).
FastEthernet0/0
P 10.37.100.0/24, 1 successors, FD is 409600 via 10.1.2.2 (409600/128256).
FastEthernet0/0
P 10.55.72.0/24. 1 successors, FD is 409600 via 10.1.2.2 (409600/128256),
FastEthernet0/0
P 10.55.144.0/24. 1 successors, FD is 409600 via 10.1.2.2 (409600/128256),
FastEthernet0/0
P 10.123.187.0/24.0 successors, FD is Inaccessible via 10.1.2.2 (409600/128256),
FastEthernet0/0
```

Refer to the exhibit. Router R2 should be learning the route for 10.123.187.0/24 via EIGRP. Which action resolves the issue without introducing more issues?

A. Redistribute the route in EIGRP with metric, delay, and reliability.

B. Use distribute-list to modify the route as an internal EIGRP route.

C. Use distribute-list to filter the external routes in OSPF.

D. Remove route redistribution in R2 for this route in OSPF.

Correct Answer: *C*

Community vote distribution

D (62%) C (38%)

ZamanR 5 days, 2 hours ago

I think D is correct

upvoted 1 times

fizzer 3 months, 2 weeks ago

C is the right answer, D is not actually possible because even if this was a RIP learned route on R2 with AD of 120, redistributing RIP into OSPF on R2 will not remove the RIP route from the routing table and install the now External OSPF route because OSPF has AD of 110.

it does makes sense when you think about it, because RIP would be the primary protocol that learned the route on R2, and redistributing it into OSPF on this same router does not make OSPF the boss over the route

Also if you look closely, OSPF learned about the route from 10.1.3.2 which is probably where the redistribution happened whereas EIGRP learned it from 10.1.2.2 who probably also did the redistribution

upvoted 2 times

chaocheng 4 months ago

ANS: C
LAB test

```
ip prefix-list 1 deny 10.123.187.0/24

router ospf 1
redistribute eigrp 1 metric 1 subnets
distribute-list prefix 1 in
```

upvoted 1 times

Cyril_the_Squirl 4 months, 1 week ago

D is correct.
OSPF has AD=110 lower than EIGRP EX=170, the prefix that makes it into the routing table is therefore OSPF.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

Sorry team I was wrong in the previous answer, analyzing the question well. The correct response is C and not D, because D is receiving the route and redistribution is not removed. The most certain thing is that the network in EIGRP has AD higher than the AD of external routes of OSPF.

Tested in lab filter external ospf route and put EIGRP route

upvoted 2 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

The option correct is D. team is logical. beacuse. there is a problem with reditribute protocol EIGRP into OSPF -> OE

upvoted 1 times

Malasxd 7 months, 1 week ago

Selected Answer: D

Definily D.

C would work, but you would impact the other external routes in OSPF

upvoted 4 times

HungarianDish 7 months, 3 weeks ago

Selected Answer: D

My assumption: the route 10.123.187.0/24 could be a static route which is redistributed into eigrp, and so it gets the AD 170 as eigrp external route.
The the route goes to R2 where all eigrp routes are redistributed into ospf as E2 external ospf routes.
At this point, the route 10.123.187.0/24 has an AD 110 in ospf and an AD 170 in eigrp on R2, thus ospf wins, and R2 learns the route from ospf.
We would need to stop the EIGRP route (10.123.187.0/24) from getting redistributed into OSPF using a route-map, which means solution "D".

upvoted 3 times

HungarianDish 7 months, 3 weeks ago

Some good examples with different solutions:
https://learningnetwork.cisco.com/s/question/0D56e0000B7yzCVCQY/filtering-of-prefix-into-out-of-both-eigrp-and-ospf
https://community.cisco.com/t5/other-network-architecture-subjects/redistribution-from-eigrp-to-ospf/td-p/290844

upvoted 1 times

Typovy 8 months ago

Selected Answer: D

D is the correct answer.
C will introduce more troubles, there are more than this one OSPF External routes so we will block all of them

upvoted 2 times

Mad_Scorpion 10 months, 2 weeks ago

Selected Answer: C

Option C verified in Lab.

upvoted 4 times

6dd4aa0 8 months, 2 weeks ago

Can we see your code?

upvoted 1 times

Patrick1234 10 months, 4 weeks ago

Since R2 is receiving the routes as "O E2" routes, he can't be the router that is redistributing them into OSPF, so D is not correct (he would see the routes as the original protocol). I think the EIGRP route to 10.123.187.0/24 is an external route and so we need to lower the AD.

I think answer B is correct:

https://community.cisco.com/t5/routing/eigrp-fd-is-inaccessible-when-re-distribution/td-p/1497303

However, i can't find a way to do this in a route map. Does anyone know if this is possible? If it's not possible, C is the only valid option, but as other have said, this will create other issues since 2 other routes will also be removed. Changing the AD of "D EX" routes would be a better option but it's not in any answer.

Also, if you do B without route filtering, it would cause new problems (loop). So the best option might actually be C here...

upvoted 1 times

Zizu007 11 months, 2 weeks ago

Selected Answer: C

answer is correct!
tested in lab.

upvoted 1 times

VergilP 1 year, 1 month ago

Selected Answer: D

I'm going for D
the question say Which action resolves the issue" without introducing more issues"
if filter the external routes in OSPF -> might delete other O E2 route??
so... I think is D

upvoted 3 times

chris7890 1 year, 2 months ago

Doesn't answer D make more sense?

upvoted 3 times

Huntkey 1 year, 2 months ago

Selected Answer: C

Googling this error got me this:
In other words, when the FD is inaccessible in the EIGRP topology table, the router is not using that EIGRP route in its routing table. Usually, the route is overridden by another routing protocol that has lower administrative distance.

upvoted 1 times

> **DUBC89x** 1 year ago
>
> C cannot be correct as it would block the other 2 external OSPF routes. Leaving D as the correct answer.
>
> upvoted 1 times

Question #220 *Topic 1*

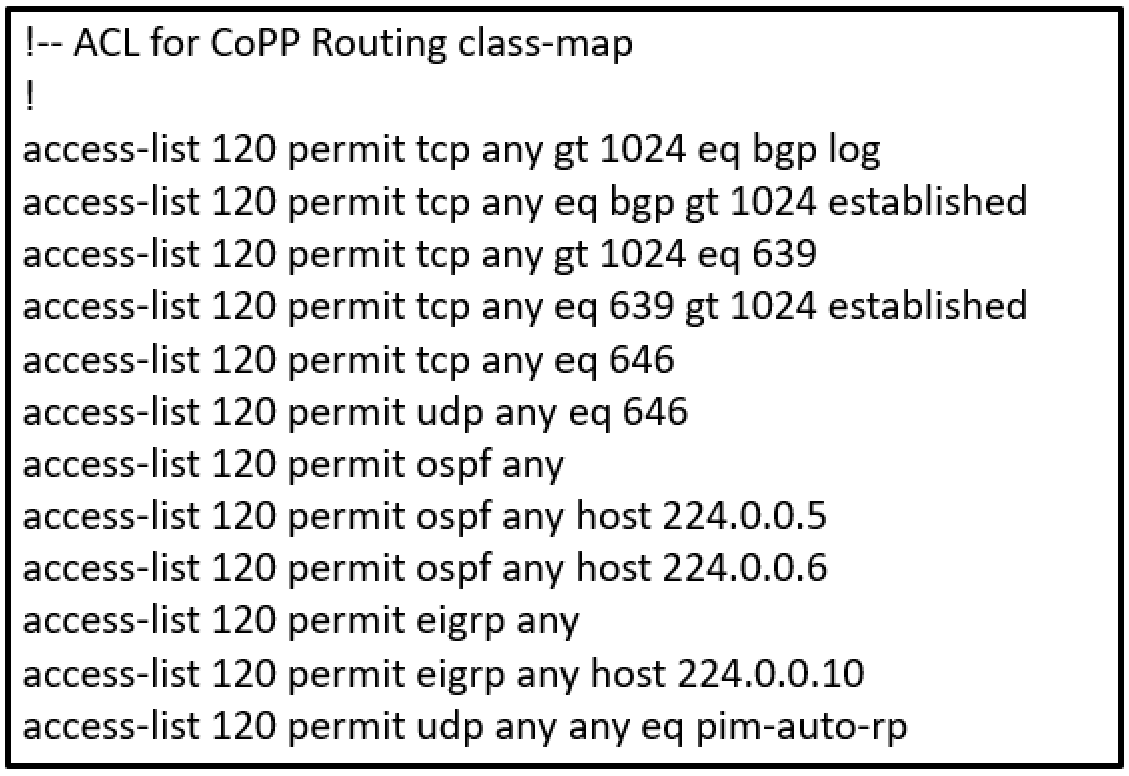

```
!-- ACL for CoPP Routing class-map
!
access-list 120 permit tcp any gt 1024 eq bgp log
access-list 120 permit tcp any eq bgp gt 1024 established
access-list 120 permit tcp any gt 1024 eq 639
access-list 120 permit tcp any eq 639 gt 1024 established
access-list 120 permit tcp any eq 646
access-list 120 permit udp any eq 646
access-list 120 permit ospf any
access-list 120 permit ospf any host 224.0.0.5
access-list 120 permit ospf any host 224.0.0.6
access-list 120 permit eigrp any
access-list 120 permit eigrp any host 224.0.0.10
access-list 120 permit udp any any eq pim-auto-rp
```

Refer to the exhibit. The control plane is heavily impacted after the CoPP configuration is applied to the router. Which command removal lessens the impact on the control plane?

A. access-list 120 permit tcp any gt 1024 eq bgp log

B. access-list 120 permit ospf any

C. access-list 120 permit udp any any eq pim-auto-rp

D. access-list 120 permit eigrp any host 224.0.0.10

Correct Answer: *A*

Community vote distribution

A (100%)

Koume Highly Voted 11 months ago

Selected Answer: A

The Real explanation here is that log option in ACL is Known to have trouble with CoPP.
"•CoPP does not support ACEs with the log keyword"
https://www.cisco.com/en/US/docs/general/Test/dwerblo/broken_guide/copp.html

upvoted 6 times

HungarianDish 7 months, 3 weeks ago

Agree, hits to ACL entry with log increase CPU utilization, because logging is done by the main CPU.

upvoted 3 times

Huntkey Most Recent 1 year, 2 months ago

I guess the only explanation is that only BGP can reach far from anywhere to the router to cause the high impact. Others like OSPF and EIGRP and LDP are from the local segment so not going to cause trouble.

upvoted 1 times

Question #221 Topic 1

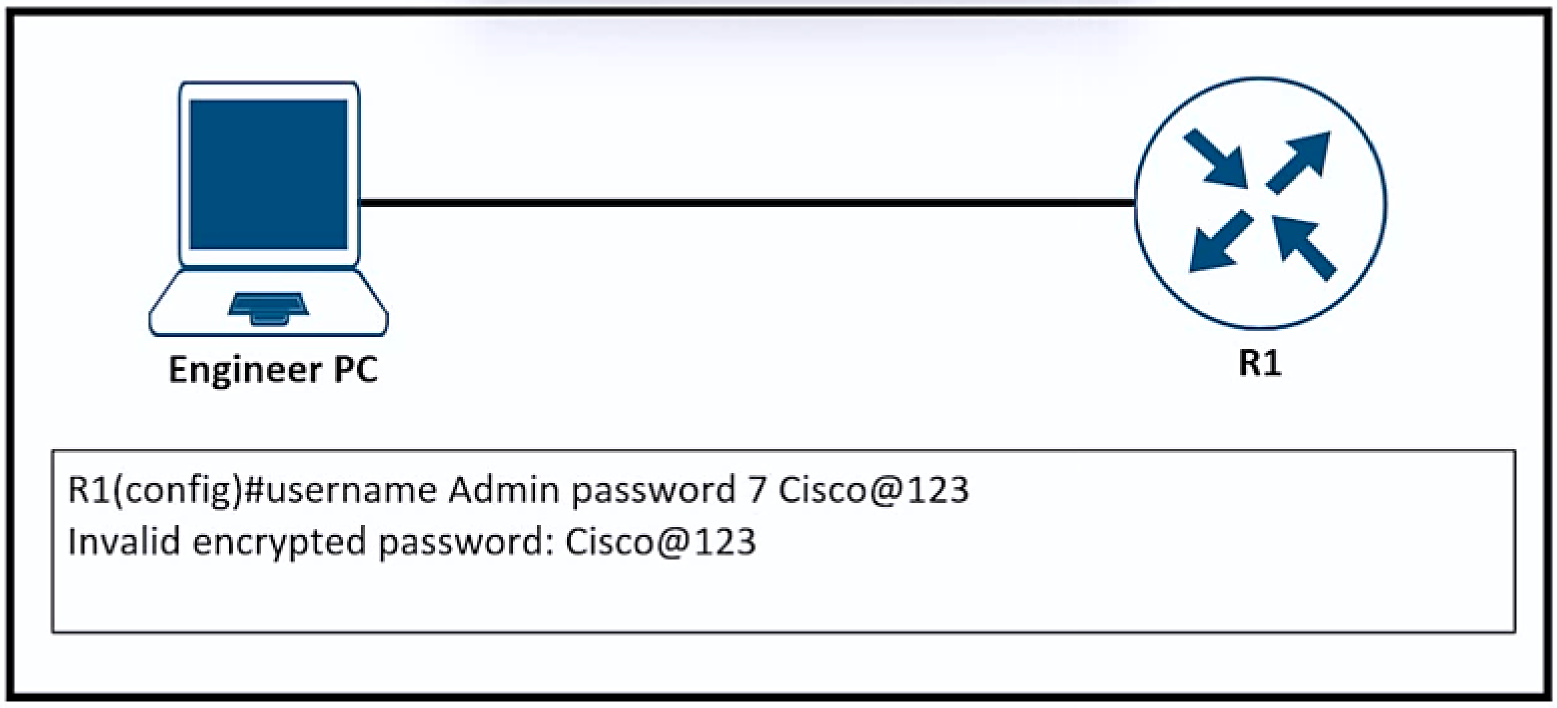

```
R1(config)#username Admin password 7 Cisco@123
Invalid encrypted password: Cisco@123
```

Refer to the exhibit. An engineer is trying to add an encrypted user password that should not be visible in the router configuration. Which two configuration commands resolve the issue?
(Choose two.)

A. username Admin password Cisco@123

B. service password-encryption

C. username Admin secret Cisco@123

D. password encryption aes

E. no service password-encryption

F. username Admin password 5 Cisco@123

Correct Answer: *BC*

Community vote distribution

AB (100%)

Brand 3 months, 4 weeks ago

Selected Answer: AB

should be A and B

upvoted 1 times

guy276465281819372 5 months ago

Selected Answer: AB

A&B. The question states ENCRYPTION not HASHING.

upvoted 1 times

MicMillon 5 months, 1 week ago

Selected Answer: AB

A|B
password will encrypt

upvoted 2 times

robi1020 5 months, 2 weeks ago

In the data security field, encryption and hashing are commonly compared, but why is this the case. Encryption is a two-way function where data is passed in as plaintext and comes out as ciphertext, which is unreadable. Since encryption is two-way, the data can be decrypted so it is readable again. Hashing, on the other hand, is one-way, meaning the plaintext is scrambled into a unique digest, through the use of a salt, that cannot be decrypted. Technically, hashing can be reversed, but the computational power needed to decrypt it makes decryption infeasible.

upvoted 2 times

guy276465281819372 5 months, 3 weeks ago

Selected Answer: AB

I believe A & B would be a suitable answer.
using "secret" HASHes the password, not encrypting it.
the engineer tried to encrypt the password not HASH it so A would be good.

upvoted 1 times

Question #222

Topic 1

A customer reports that traffic is not passing on an EIGRP enabled multipoint interface on a router configured as below:

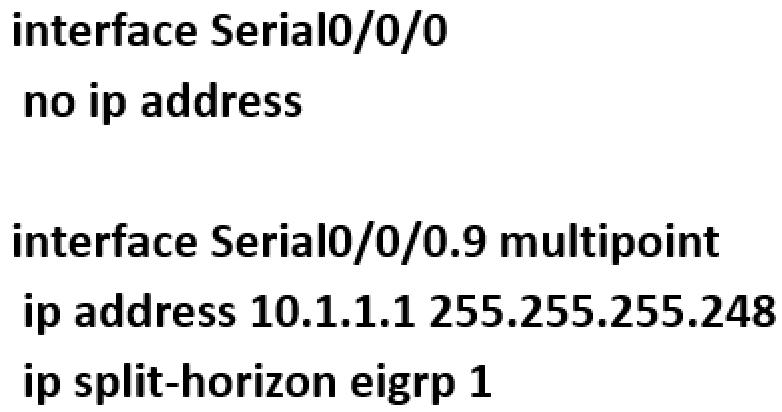

```
interface Serial0/0/0
 no ip address

interface Serial0/0/0.9 multipoint
 ip address 10.1.1.1 255.255.255.248
 ip split-horizon eigrp 1
```

Which action resolves the issue?

A. Enable poison reverse.

B. Disable split horizon.

C. Disable poison reverse.

D. Enable split horizon.

Correct Answer: *B*

Community vote distribution

B (100%)

MasterMatt 7 months, 4 weeks ago

Selected Answer: B

Poison Reverse is used in RIP enabled interfaces to tackle the count-to-Infinity problems and one can imagine it as a reverse of the Split Horizon method.

Split horizon does not send routes learned from the same interface to avoid loops.

upvoted 1 times

GodFather 10 months, 3 weeks ago

When split horizon is enabled, any route learned from an interface is not advertised back out the same interface. This rule is intended to stop routing loops with distance-vector protocols.

upvoted 2 times

Question #223 Topic 1

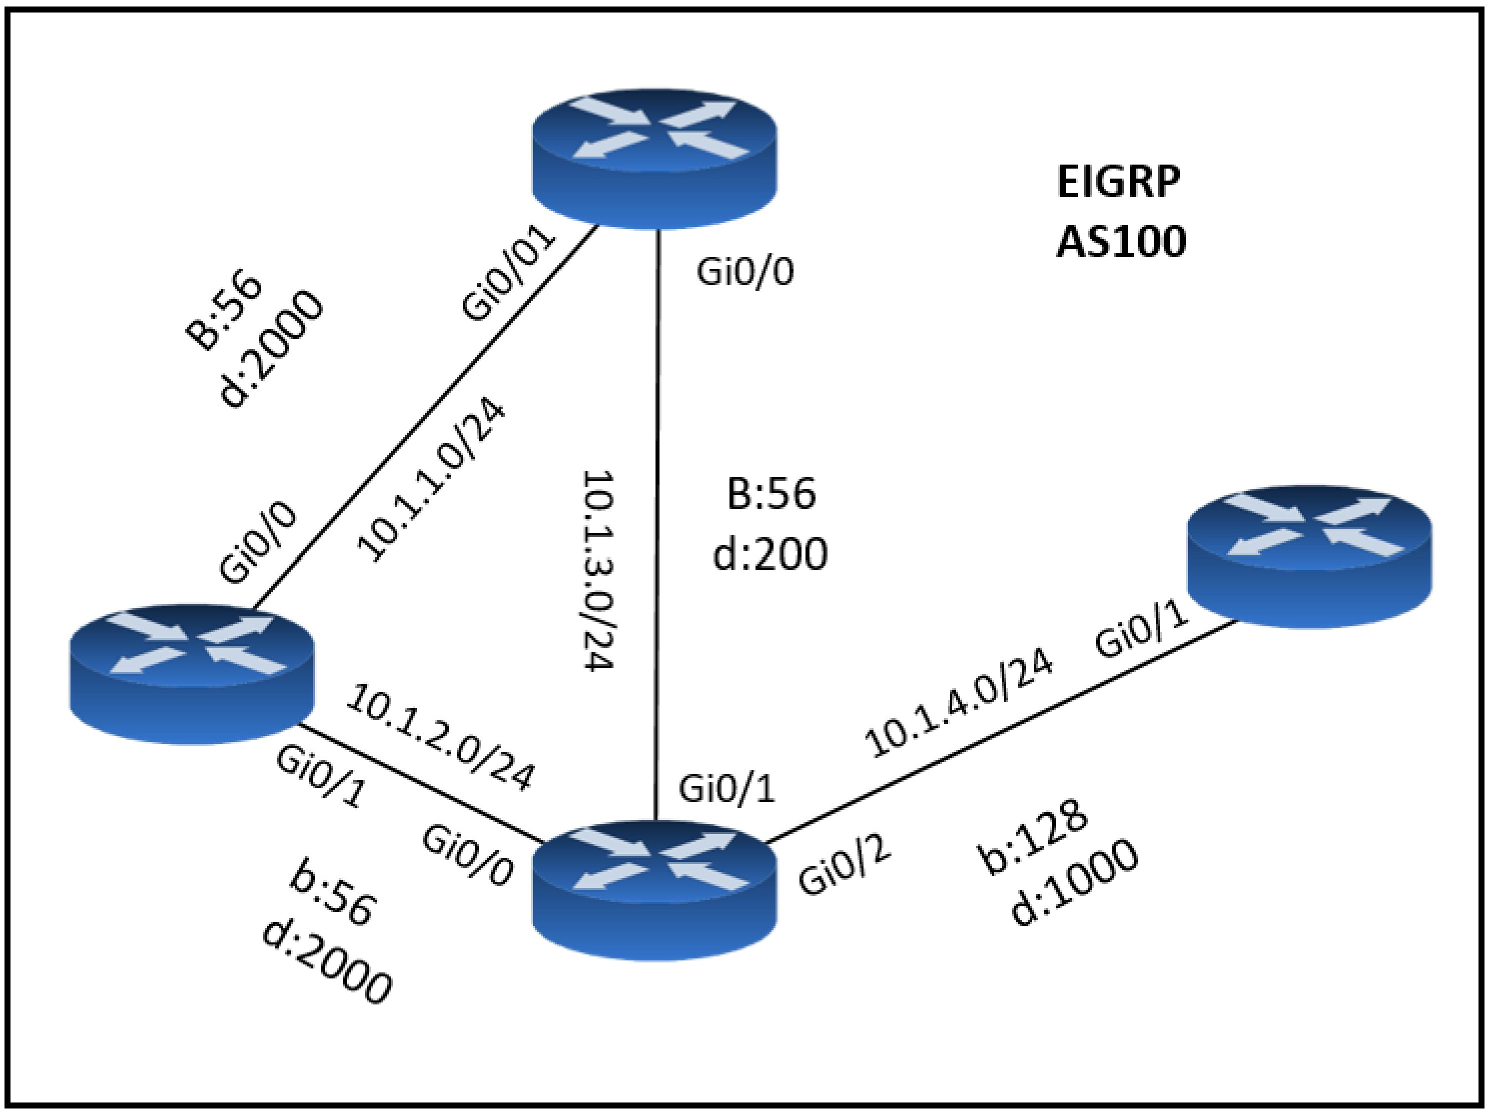

Refer to the exhibit. A loop occurs between R1, R2, and R3 while EIGRP is run with poison reverse enabled. Which action prevents the loop between R1, R2, and
R3?

A. Enable split horizon.

B. Configure R3 as stub receive-only.

C. Configure route tagging.

D. Configure route filtering.

Correct Answer: *A*

Community vote distribution

A (56%) D (44%)

ZamanR 5 days, 5 hours ago

A is the correct answer

upvoted 1 times

SAMAKEMM 2 months, 2 weeks ago

Selected Answer: A

Split harizon is enable to prevent loop in EIGRP

upvoted 2 times

chris110 3 months, 1 week ago

Selected Answer: D

In Cisco devices, split horizon is always used along with poison reverse (via the command “ip split-horizon”) so in this question split horizon is already turned on. To prevent loop we can only use route filtering.

upvoted 2 times

diegodavid82 4 months ago

Selected Answer: A

It's the correct answer, review the document provided by HungarianDish. Both Split horizon and poison reverse works together for resolve this issue

upvoted 3 times

inteldarvid 5 months, 1 week ago

https://notes.networklessons.com/eigrp-split-horizon-vs-poison-reverse

upvoted 2 times

HamzaBadar 5 months, 3 weeks ago

Selected Answer: D

Split horizon is always used with poison reverse in cisco devices. therefore, the only solution is route filtering.

upvoted 2 times

Malasxd 7 months, 2 weeks ago

I thinks it's "A". Split Horizon, Poison reverse and feaseble condition are the mechanisms EIGRP uses to prevents loops.

upvoted 2 times

HungarianDish 7 months, 3 weeks ago

EIGRP combines poison reverse and split horizon to help prevent routing loops. So, if the question is seeking some general answer, then probably it is "A". Further information (about the loop or about the design) is required to give an accurate answer.

upvoted 2 times

HungarianDish 7 months, 3 weeks ago

The question is based on this cisco document:
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html#anc21

upvoted 2 times

HungarianDish 7 months, 3 weeks ago

B) stub receive-only -> I would not consider it as a loop prevention mechanism, so for me this answer is excluded. It prevents Stuck In Active.
https://networklessons.com/eigrp/eigrp-stub-explained
https://networklessons.com/eigrp/eigrp-queries-and-stuck-in-active
https://www.geeksforgeeks.org/configuring-eigrp-stub-in-cisco/

upvoted 1 times

HungarianDish 7 months, 3 weeks ago

Topology can be viewed in other dumps:
https://vceguide.com/which-action-prevents-the-loop-between-r1-r2-and-r3/

upvoted 2 times

Dacusai 7 months, 3 weeks ago

First who is R1, R2, R3 and R4, second split horizon has nothing to do here because routers are not sending routes back from int. it was learned. I think B is the correct one on this case.

upvoted 1 times

> **Malasxd** 7 months, 2 weeks ago
>
> The EIGRP routers exchange full routing table with each other. They don't send routes back because the split horizon. So it has a lot to here
>
> upvoted 1 times

chris7890 11 months, 2 weeks ago

i think the given answer is correct.

upvoted 1 times

Huntkey 1 year, 2 months ago

I guess that the answer is correct. Only by disabling the split-horizon could cause a loop in an all EIGRP topology

upvoted 1 times

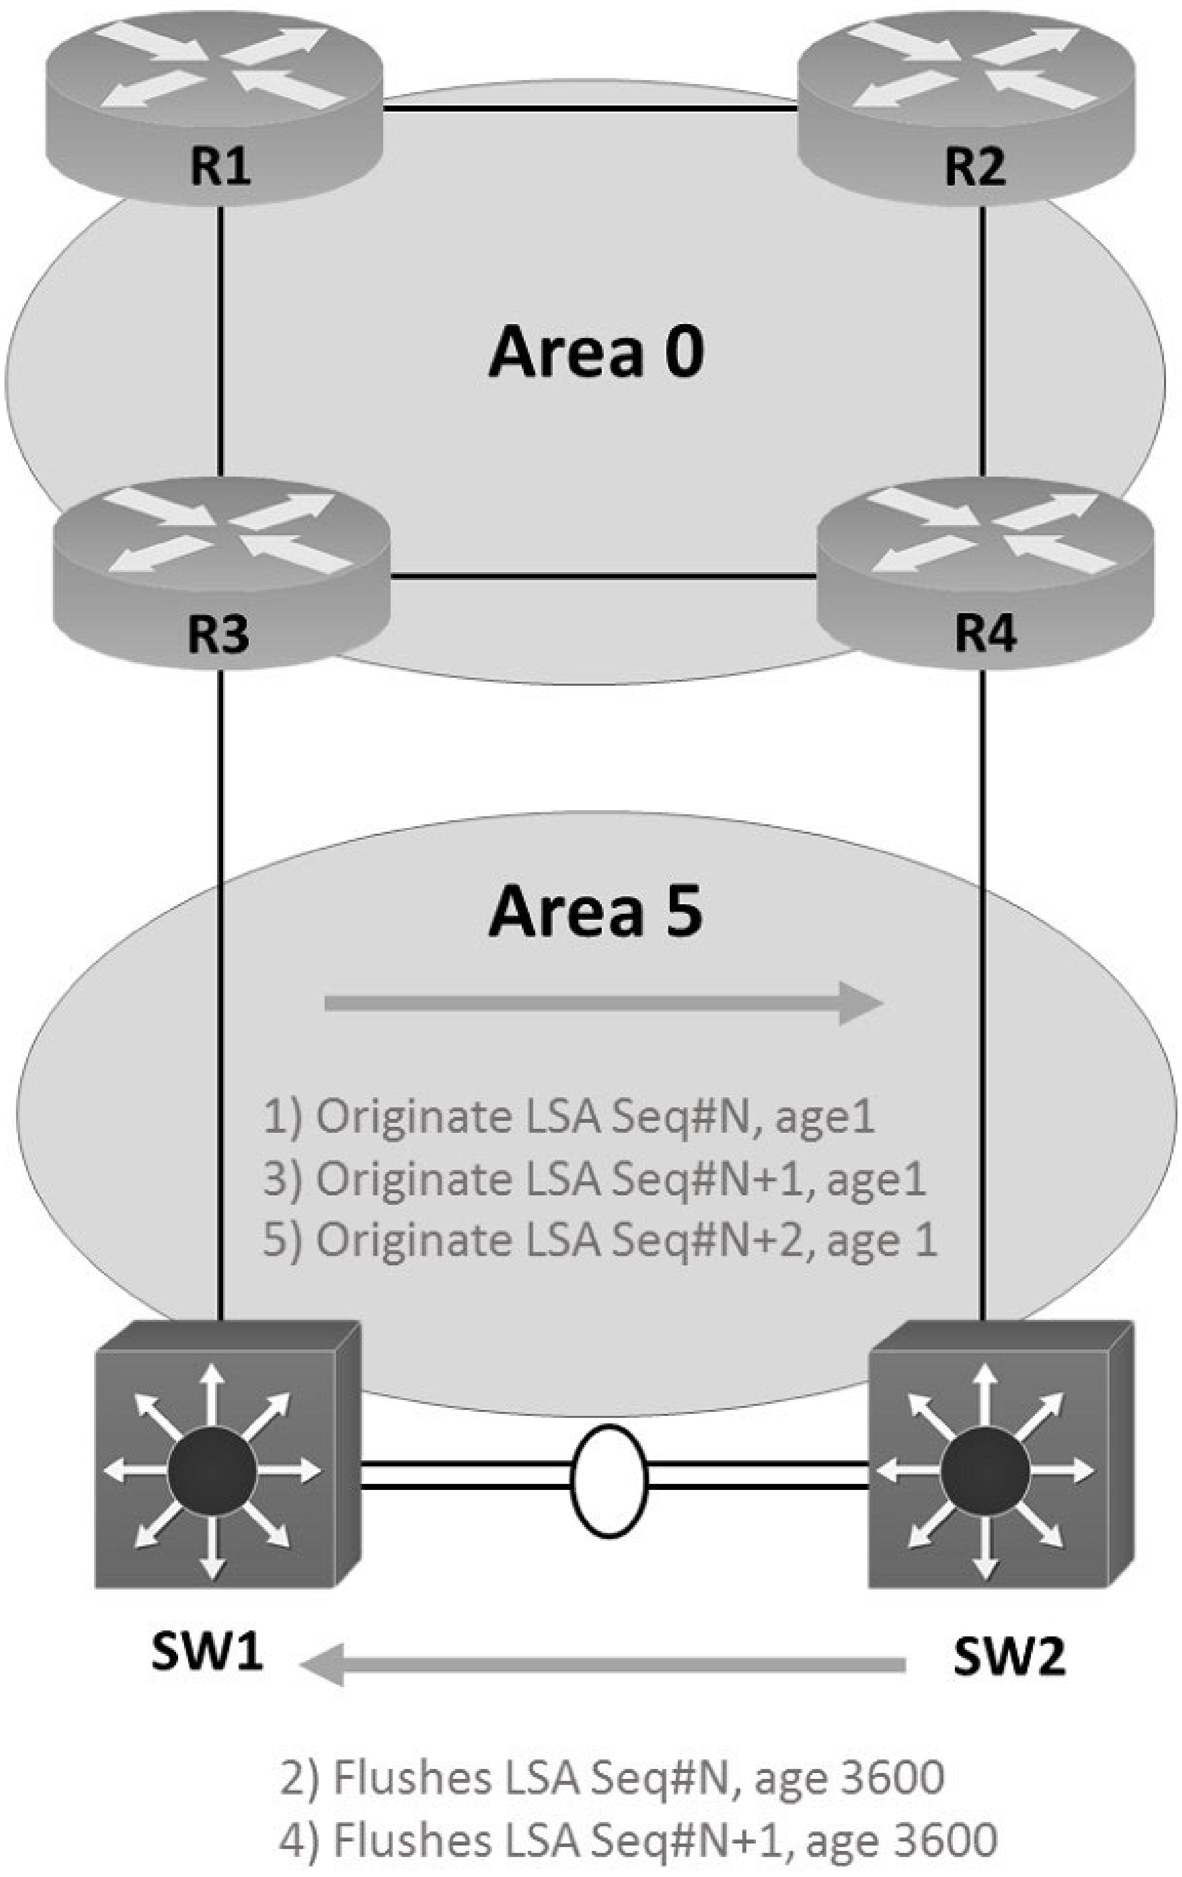
R1
R2
R3
R4
Area 0
Area 5
1) Originate LSA Seq#N, age1
3) Originate LSA Seq#N+1, age1
5) Originate LSA Seq#N+2, age 1
SW1
SW2
2) Flushes LSA Seq#N, age 3600
4) Flushes LSA Seq#N+1, age 3600

Refer to the exhibit. An error message "an OSPF-4-FLOOD_WAR" is received on SW2 from SW1. SW2 is repeatedly receiving its own link-state advertisement and flushes it from the network. Which action resolves the issue?

A. Change area 5 to a normal area from a nonstub area.

B. Resolve different subnet mask issue on the link.

C. Configure Layer 3 port channel on interfaces between switches.

D. Resolve duplicate IP address issue in the network.

Correct Answer: *D*

Community vote distribution

D (100%)

Stylar Highly Voted 6 months ago

More questions like this and i will give up

upvoted 6 times

inteldarvid Most Recent 5 months, 1 week ago

Selected Answer: D

team is option D:
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/118880-technote-ospf-00.html

upvoted 1 times

HungarianDish 7 months, 3 weeks ago

Selected Answer: D

The question is based on this cisco document:
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/118880-technote-ospf-00.html

"This is meant to detect issues with Type-2 LSAs when duplicate IP addresses are present in the network, or with Type-5 LSAs when there is a duplicate router ID in different OSPF Areas."

upvoted 3 times

Xerath 9 months, 4 weeks ago

Selected Answer: D

https://community.cisco.com/t5/switching/ospf-4-flood-war-messages-after-config-change/td-p/2506500

upvoted 1 times

NoUserName1234 1 year, 3 months ago

Seems right see:
http://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/9237-9.html

upvoted 1 times

Question #225 *Topic 1*

Which two components are required for MPLS Layer 3 VPN configuration? (Choose two.)

A. Use LDP for customer routes.

B. Use pseudowire for Layer 2 routes.

C. Use a unique RD per customer VRF.

D. Use OSPF between PE and CE.

E. Use MP-BGP for customer routes.

Correct Answer: *CD*

Community vote distribution

CE (100%)

HungarianDish 7 months, 3 weeks ago

Selected Answer: CE

https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2020/pdf/DGTL-BRKMPL-1100.pdf

upvoted 2 times

elmones 8 months, 1 week ago

But MP-BGP does not propagate the customer routes, propagate VPNv4 routes

upvoted 1 times

heeeeyajoke 1 year ago

Definitely C and E

upvoted 1 times

mrnipsnips 1 year, 1 month ago

Selected Answer: CE

CE like the others explained

upvoted 1 times

Remsync 1 year, 2 months ago

Selected Answer: CE

You need a RD and MP-BGP to propagate the customer routes through the MPLS. OSPF between the PE and CE *can* be used and will work fine, but is not needed.

CE

upvoted 2 times

Huntkey 1 year, 2 months ago

Selected Answer: CE

It can be any routing protocol between CE and PE including static routes

upvoted 3 times

Question #226 Topic 1

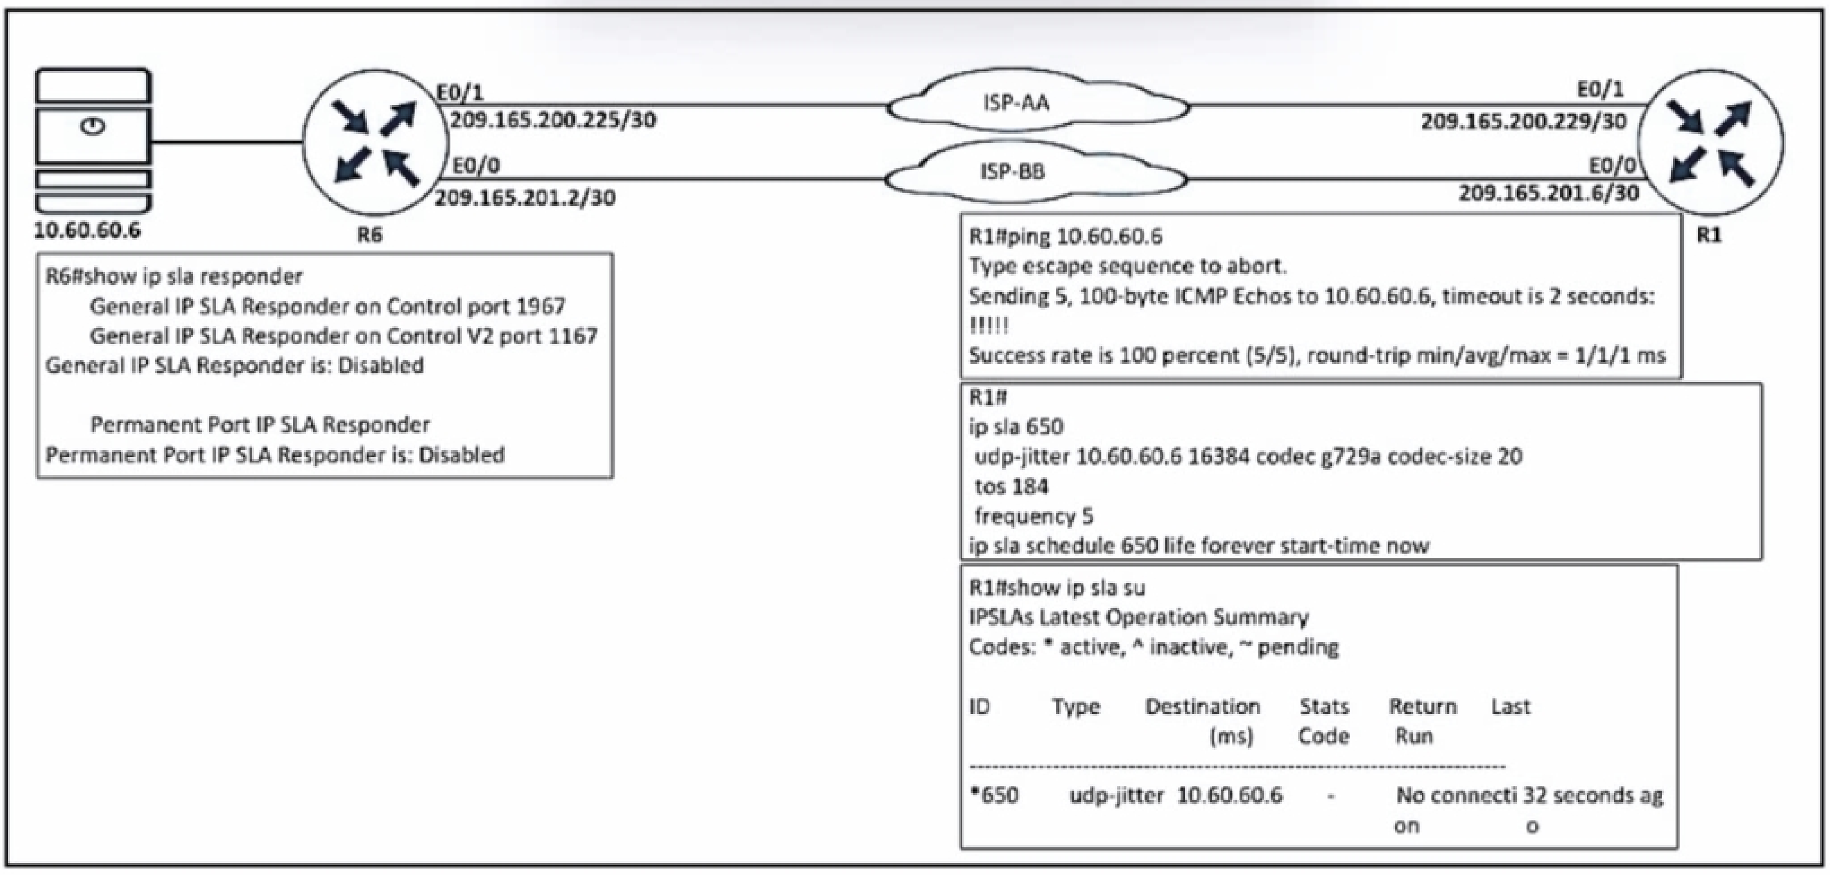

```
R6#show ip sla responder
     General IP SLA Responder on Control port 1967
     General IP SLA Responder on Control V2 port 1167
General IP SLA Responder is: Disabled

     Permanent Port IP SLA Responder
Permanent Port IP SLA Responder is: Disabled
```

```
R1#ping 10.60.60.6
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.60.60.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

```
R1#
ip sla 650
 udp-jitter 10.60.60.6 16384 codec g729a codec-size 20
 tos 184
 frequency 5
ip sla schedule 650 life forever start-time now
```

```
R1#show ip sla su
IPSLAs Latest Operation Summary
Codes: * active, ^ inactive, ~ pending

ID       Type       Destination     Stats     Return     Last
                          (ms)      Code       Run
---------------------------------------------------------------
*650      udp-jitter  10.60.60.6      -       No connecti 32 seconds ag
                                              on               o
```

Refer to the exhibit. Which configuration resolves the IP SLA issue from R1 to the server?

A. R6(config)#ip sla responder

B. R6(config)#ip sla 650 R6(config-ip-sla)#udp-jitter 10.60.60.6

C. R6(config)#ip sla responder udp-echo ipaddress 10.60.60.6 po 5000

D. R6(config)#ip sla schedule 10 life forever start-time now

Correct Answer: *C*

Community vote distribution

A (92%) 8%

AlexInShort12 4 days ago

Selected Answer: A

One point I should mention is that the IP SLA Responder is not required for IP SLA to function, but it does allow for more detailed information gathering and reporting.
https://learningnetwork.cisco.com/s/blogs/a0D3i000002SKN0EAO/ip-sla-fundamentals

upvoted 1 times

mouin 3 months, 1 week ago

Selected Answer: A

I tested option A and option C
option A is correct

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: A

option A , i test in my lab

upvoted 1 times

HungarianDish 7 months, 3 weeks ago

Selected Answer: A

Sorry for previous post, the ip address in answer "C" is incorrect. An ip address on R1 should be the destination IP for the reflected UDP traffic.

upvoted 3 times

HungarianDish 7 months, 3 weeks ago

Selected Answer: C

Based on the output, they want to measure udp-jitter. To configure IP SLA responders for UDP jitter use:
#ip sla responder udp-echo ipaddress <> port <>

upvoted 1 times

Commando1664 8 months ago

I just labbed it and the out come is C

upvoted 2 times

Zizu007 11 months, 2 weeks ago

Selected Answer: A

```
R5#show ip sl summary
IPSLAs Latest Operation Summary
Codes: * active, ^ inactive, ~ pending

ID Type Destination Stats Return Last
(ms) Code Run
-----------------------------------------------------------------------
*5 udp-jitter 6.6.6.6 RTT=6 OK 0 seconds ago
----
LOCAL:
ip sla 5
udp-jitter 6.6.6.6 5000
threshold 1000
timeout 2000
frequency 2
ip sla schedule 5 life forever start-time now
-----
REMOTE:
ip sla responder
```

upvoted 3 times

heeeeyajoke 1 year ago

i think you only need one line of command to enable ip sla responder. A for me

upvoted 1 times

mrnipsnips 1 year, 1 month ago

Selected Answer: A

I'm voting A, too lazy to lab it tho haha

upvoted 1 times

Huntkey 1 year, 2 months ago

Selected Answer: A

C doesn't work at all for some reason. I did notice that it is udp-jitter on one side and udp-echo on the responder side. However, I tried with the udp-echo as well and the C alone still doesn't work. It must be A then

upvoted 1 times

jarz 1 year, 2 months ago

Selected Answer: A

```
# ip sla responder
```

upvoted 1 times

Question #227 *Topic 1*

A network administrator added a new spoke site with dynamic IP on the DMVPN network. Which configuration command passes traffic on the DMVPN tunnel from the spoke router?

A. ip nhrp registration no-registration

B. ip nhrp registration dynamic

C. ip nhrp registration no-unique

D. ip nhrp registration ignore

Correct Answer: *C*

Community vote distribution

C (100%)

inteldarvid 5 months, 2 weeks ago

Selected Answer: C

option C is correct

https://community.cisco.com/t5/security-knowledge-base/unable-to-pass-traffic-on-the-dynamic-multipoint-vpn-tunnel-with/ta-p/3111776#:~:text=If%20you%20configure%20the%20ip,such%20as%20a%20dial%20environment.

upvoted 1 times

jarz 1 year, 2 months ago

Non-Unique Registrations
If you’re experiencing DMVPN downtime due to changing public IP addresses of your DMVPN spokes, apply the ip nhrp registration non-unique interface configuration command to the DMVPN tunnel interface. This command will reduce the recovery time to less than a minute. Faster recovery is harder to achieve as the router has to execute a number of steps following a physical interface flap:

Install new static routes to the hub sites;
Create IPSec session with the hub sites;
Register new public IP address with NHRP;
Establish routing adjacency.
You can fine-tune steps 1-3 on the spoke router; step 4 sometime requires coordinated changes throughout the network.

https://blog.ipspace.net/2010/09/dmvpn-non-unique-nhrp-registrations.html

upvoted 3 times

NoUserName1234 1 year, 3 months ago

seems right:
https://yurmagccie.wordpress.com/2016/06/07/dmvpn/
Search for 'no-unique'

upvoted 1 times

Question #228 *Topic 1*

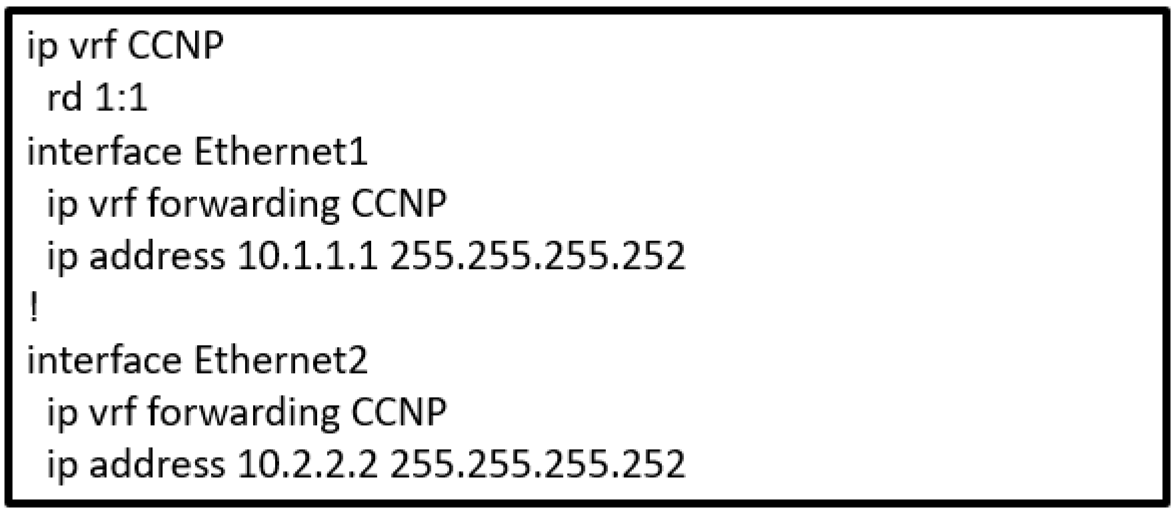

```
ip vrf CCNP
  rd 1:1
interface Ethernet1
  ip vrf forwarding CCNP
  ip address 10.1.1.1 255.255.255.252
!
interface Ethernet2
  ip vrf forwarding CCNP
  ip address 10.2.2.2 255.255.255.252
```

Refer to the exhibit. Which configuration enables OSPF for area 0 interfaces to establish adjacency with a neighboring router with the same VRF?

A. router ospf 1 vrf CCNP network 10.1.1.1 0.0.0.0 area 0 network 10.2.2.2 0.0.0.0 area 0

B. router ospf 1 interface Ethernet1 ip ospf 1 area 0.0.0.0 interface Ethernet2 ip ospf 1 area 0.0.0.0

C. router ospf 1 vrf CCNP interface Ethernet1 ip ospf 1 area 0.0.0.0 interface Ethernet2 ip ospf 1 area 0.0.0.0

D. router ospf 1 vrf CCNP network 10.0.0.0 0.0.255.255 area 0

Correct Answer: *C*

Community vote distribution

C (54%) A (46%)

Fenix7 3 months, 1 week ago

Can't be C because the area is 0, and not 0.0.0.0. The answer is A

upvoted 1 times

fizzer 3 months, 1 week ago

Option C seems like the best bet seeing as "Interface" was stressed in the question

Both configuration works as already highlighted by others, however, "show ip protocols" shows option A's configuration under "Routing for networks" whereas it shows option C's configuration under "Routing for Interfaces configured explicitly for Area:"

I think the idea behind the question is which of the 2 configuration commands put the interface under the explicit configuration in "show ip protocols"

Option A is intentionally meant to sway, because it uses the Interface IP address rather than the network address, however it does not show under explicit interface configuration in "show ip protocols"

upvoted 1 times

Wolfxx 4 months ago

I agree with answer "C", because when question says "Which configuration enables OSPF for area 0 interfaces", it's closer.

upvoted 1 times

Malasxd 7 months ago

Selected Answer: C

"A" and "C" works.
I choose "C" due to keyword "interfaces". The chance to be wrong is lower hehehe

upvoted 3 times

HungarianDish 7 months, 3 weeks ago

I also labbed it in CML. Same result as for Huntkey. Both "A" and "C" work. "A" uses 0.0.0.0 wildcard masks in the network statement, so ospf is enabled only on a specific interface. "C" is associating the ospf process directly under the interface configuration. Both solutions seems to be OK.

upvoted 2 times

azzawim 8 months ago

Selected Answer: C

Question mention interface

upvoted 3 times

forccnp 10 months, 1 week ago

Selected Answer: C

Key word in the question is 'interfaces', C is the correct one

upvoted 1 times

ttt00909 11 months, 2 weeks ago

Selected Answer: A

A desu

upvoted 2 times

PimplePooper 12 months ago

Selected Answer: A

A is correct. Both interfaces fall within the ospf network statements.

upvoted 3 times

jarz 1 year, 1 month ago

I'm leaning toward C as both A and C are valid configs. I think the key word in the question is 'interfaces', SOPF needs to enabled on interfaces only.

upvoted 3 times

> **Slinky** 8 months, 2 weeks ago
>
> I would tend to agree, but the network statements in A use 0.0.0.0 wildcard masks and thus can only apply to the IP addresses of the interfaces themselves. I suppose you could take it a step further and say that if you changes the IP on the interface then the network statement wouldn't apply anymore, but that seems unlikely. I don't love this question.
>
> upvoted 1 times

NoUserName1234 1 year, 2 months ago

Selected Answer: A

A is correct

upvoted 1 times

Huntkey 1 year, 2 months ago

I tried in the lab and both A and C work. Anything I am missing here?

upvoted 2 times

lisanta12 1 year, 3 months ago

A is answer

upvoted 3 times

Question #229 *Topic 1*

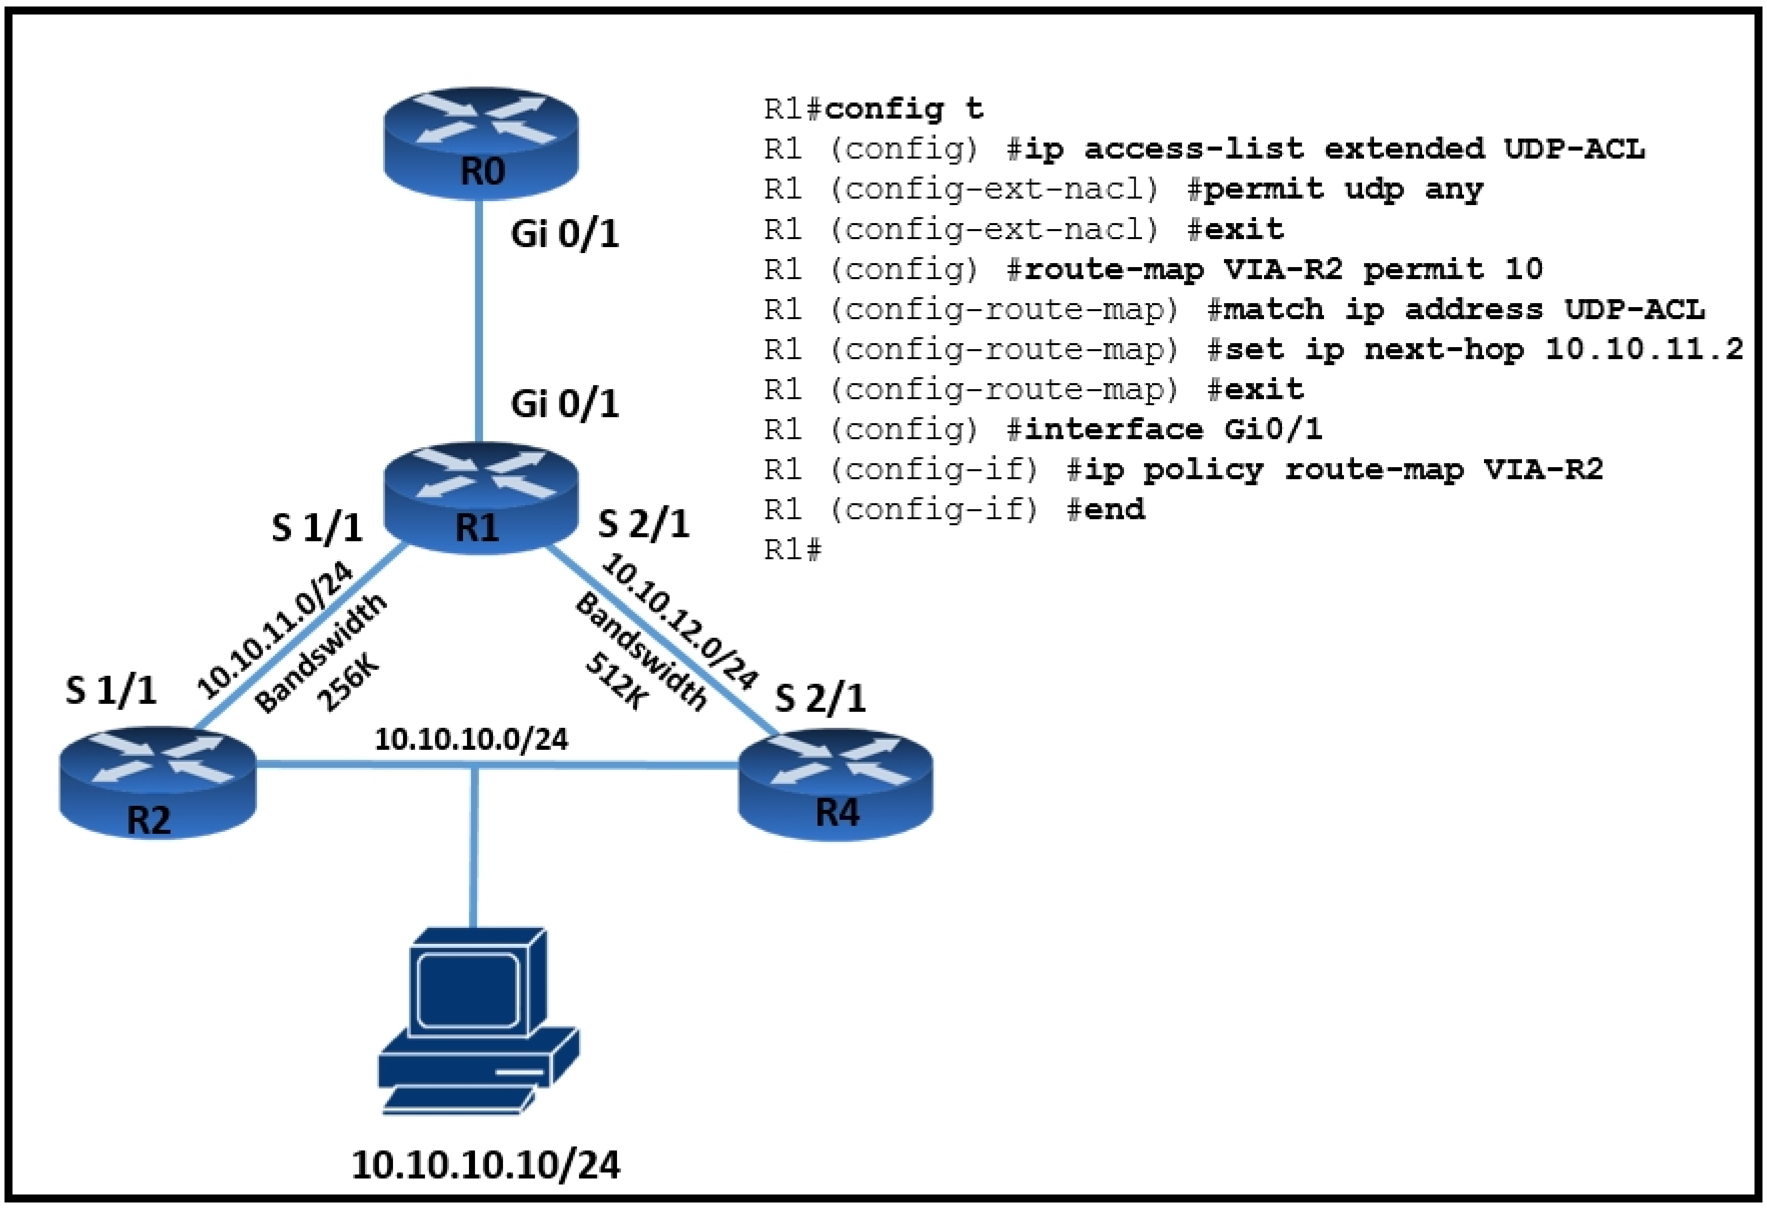

```
R1#config t
R1 (config) #ip access-list extended UDP-ACL
R1 (config-ext-nacl) #permit udp any
R1 (config-ext-nacl) #exit
R1 (config) #route-map VIA-R2 permit 10
R1 (config-route-map) #match ip address UDP-ACL
R1 (config-route-map) #set ip next-hop 10.10.11.2
R1 (config-route-map) #exit
R1 (config) #interface Gi0/1
R1 (config-if) #ip policy route-map VIA-R2
R1 (config-if) #end
R1#
```

Refer to the exhibit. TCP traffic should be reaching host 10.10.10.10/24 via R2. Which action resolves the issue?

A. Allow TCP in the access list with no changes to the route map.

B. Add a permit 20 statement in the route map to allow TCP traffic.

C. TCP traffic will reach the destination via R2 without any changes.

D. Set IP next-hop to 10.10.12.2 under the route-map permit 10 to allow TCP traffic.

Correct Answer: *A*

Community vote distribution

A (75%) B (25%)

ZamanR 5 days, 22 hours ago

A is Correct

upvoted 1 times

DeWalt95 1 week, 5 days ago

We don't know if the router are using an IGP to account for bandwidth differences..A is the best answer

upvoted 2 times

Mishranihal737 2 months, 3 weeks ago

Selected Answer: B

Why B is not feasible?

upvoted 1 times

Pietjeplukgeluk 1 week, 4 days ago

I actually agree, the question can be more to the point, i thought B initially also. A and B will solve the issue. Maybe a small reason why B could less correct: the actual placement of the permit statement is not relevant. Only TCP should be added as a permit statement, so A can be seen as more accurate here perhaps. Also A states not changing the route-map, that is correct also. So i can live with A as the correct answer.

upvoted 2 times

RamazanLokov 3 months, 1 week ago

Selected Answer: A

A is correct

upvoted 3 times

NoUserName1234 1 year, 3 months ago

A is correct, the ACL is for UDP only BUT is already forced through R2 to reach 10.10.10.10.
The BW is higher via R4 so that's the path TCP will take, so adding the TCP to the ACL forces it also via R2.

upvoted 3 times

Question #230 *Topic 1*

A newly installed spoke router is configured for DMVPN with the ip mtu 1400 command. Which configuration allows the spoke to use fragmentation with the maximum negotiated TCP MTU over GRE?

A. ip tcp adjust-mss 1360 crypto ipsec fragmentation mtu-discovery

B. ip tcp adjust-mss 1360 crypto ipsec fragmentation after-encryption

C. ip tcp payload-mtu 1360 crypto ipsec fragmentation after-encryption

D. ip tcp payload-mtu 1360 crypto ipsec fragmentation mtu-discovery

Correct Answer: *B*

Community vote distribution

B (100%)

HungarianDish 7 months, 3 weeks ago

Selected Answer: B

"When IPsec is being used, it is customary to set the MTU size on the tunnel interfaces to 1,400 bytes and to set the TCP-MSS-adjust to 1,360 bytes. "
https://www.networkworld.com/article/2224654/mtu-size-issues.html
For me "# crypto ipsec fragmentation before-encryption" would make more sense, but that option is not give, so I vote for "B".
" If the routers are performing fragmentation on behalf of the source node, it may be desirable to have the fragmentation performed prior to encryption, so the destination tunnel router doesn't have to reassemble the fragments and then perform the decryption."
https://manualzz.com/doc/33447188/configuring-ipsec-vpn-fragmentation-and-mtu

upvoted 2 times

Zizu007 11 months, 2 weeks ago

Selected Answer: B

Correct!

```
R5(config-if)#crypto ipsec fragmentation ?
after-encryption Perform fragmentation of large packets after IPSec
encapsulation
before-encryption Perform fragmentation of large packets before IPSec
encapsulation.
```

upvoted 1 times

Question #231 Topic 1

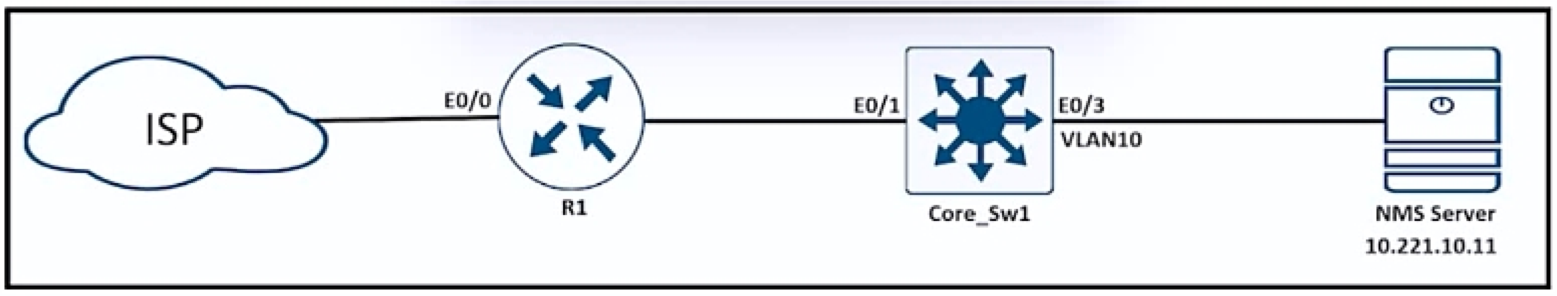

Refer to the exhibit. During ISP router maintenance, the network produced many alerts because of the flapping interface. Which configuration on R1 resolves the issue?

A. ip verify drop-rate notify hold-down 60

B. snmp trap link-status down

C. snmp trap ip verify drop-rate

D. no snmp trap link-status

Correct Answer: *D*

Community vote distribution

D (100%)

conft 4 months, 1 week ago

Selected Answer: D

The given answer is correct!

upvoted 1 times

guy276465281819372 5 months ago

Answer is correct but such a weird question, I don't think the interface being monitored is a problem.

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: D

Option correct is D: Because, D is necesary while execute maintanance windows. The option A is wrong, because that command is for uRPF: muRPF is a security feature that helps limit or even eliminate spoofed IP packets on a network.
This is accomplished by examining the source IP address of an ingress packet and determining whether it is valid. If it is valid, the packet will be forwarded. If it is not valid, the packet
Chapter 22: Infrastructure Security 853
will be discarded. Note that CEF (Cisco Express Forwarding) must be enabled on the IOS
device for uRPF to work

upvoted 2 times

ellen_AA 11 months, 1 week ago

Selected Answer: D

Given answer is correct!
https://www.oreilly.com/library/view/cisco-ios-in/0596008694/re785.html

upvoted 3 times

HungarianDish 7 months, 2 weeks ago

The answer and and the link seem to be appropriate.

upvoted 1 times

DUBC89x 1 year ago

Example:
Router(config)# ip verify drop-rate notify hold-down 60
Configures the minimum time, in seconds, between Unicast RPF drop-rate notifications.
The range is from 30 to 300. The default is 300.
"https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_urpf/configuration/xe-3s/sec-data-urpf-xe-3s-book/sec-urpf-mib-xe-3s.html"

upvoted 1 times

Question #232 Topic 1

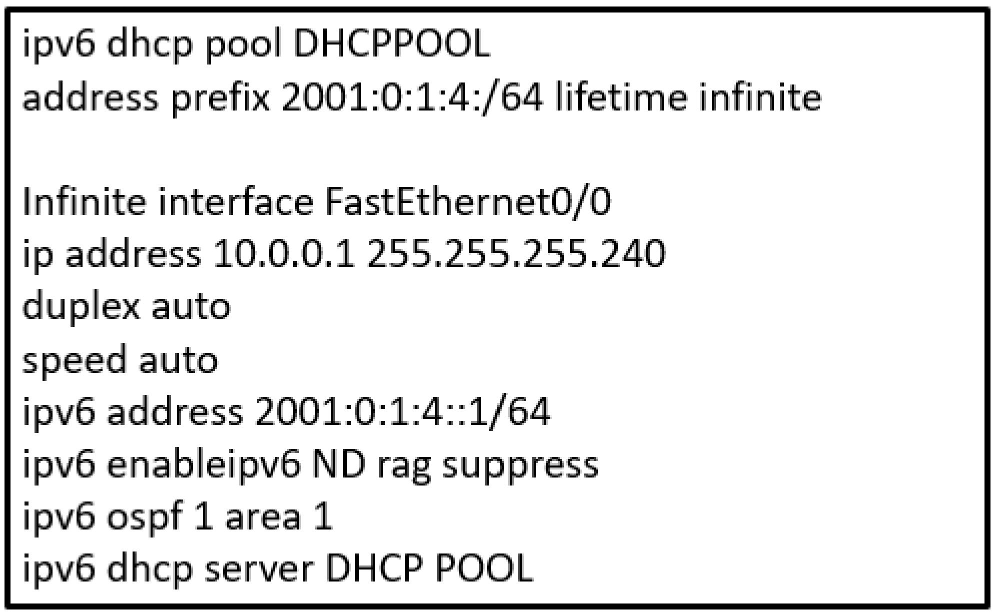

```
ipv6 dhcp pool DHCPPOOL
address prefix 2001:0:1:4:/64 lifetime infinite

Infinite interface FastEthernet0/0
ip address 10.0.0.1 255.255.255.240
duplex auto
speed auto
ipv6 address 2001:0:1:4::1/64
ipv6 enableipv6 ND rag suppress
ipv6 ospf 1 area 1
ipv6 dhcp server DHCP POOL
```

Refer to the exhibit. Reachability between servers in a network deployed with DHCPv6 is unstable. Which command must be removed from the configuration to make DHCPv6 function?

A. ipv6 nd ra suppress

B. address prefix 2001:0:1:4::/64 lifetime infinite infinite

C. ipv6 dhcp server DHCP POOL

D. ipv6 address 2001:0:1:4::1/64

Correct Answer: *A*

Community vote distribution

A (100%)

Huntkey Highly Voted 1 year, 2 months ago

Selected Answer: A

In IPv6, hosts locate a router through Router Advertisement (RA) messages sent from routers instead of by DHCP; IPv6-enabled routers that support dynamic address assignment are expected to announce themselves on the network to all clients. As such, DHCPv6 does not include any gateway information

upvoted 5 times

conft Most Recent 4 months, 1 week ago

Selected Answer: A

A is correct

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: A

OPTION A CORERCT:
https://www.networkacademy.io/ccna/ipv6/stateful-dhcpv6#:~:text=Configuring%20a%20Cisco%20router%20as%20a%20Stateful%20DHCPv6%20server&text=We%20must%20create%20a%20new,servers%2C%20and%20a%20domain%20name.

upvoted 1 times

Question #233 *Topic 1*

A customer requested a GRE tunnel through the provider network between two customer sites using loopback to hide internal networks. Which configuration on
R2 establishes the tunnel with R1?

A. R2(config)#interface Tunnel1 R2(config-if)#ip address 172.20.1.2 255.255.255.0 R2(config-if)#ip mtu 1400 R2(config-if)#ip tcp adjust-mss 1360 R2(config-if)#tunnel source 192.168.20.1 R2(config-if)#tunnel destination 192.168.10.1

B. R2(config)#interface Tunnel1 R2(config-if#ip address 172.20.1.2 255.255.255.0 R2(config-if)#ip mtu 1400 R2(config-if)#ip tcp adjust-mss 1360 R2(config-if)#tunnel source 10.10.2.2 R2(config-if)#tunnel destination 10.10.1.1

C. R2(config)#interface Tunnel1 R2(config-if)#ip address 172.20.1.2 255.255.255.0 R2(config-if)#ip mtu 1500 R2(config-if)#ip tcp adjust-mss 1360 R2(config-if)#tunnel source 10.10.2.2 R2(config-if)#tunnel destination 10.10.1.1

D. R2(config)#interface Tunnel1 R2(config-if)#ip address 172.20.1.2 255.255.255.0 R2(config-if)#ip mtu 1500 R2(config-if)#ip tcp adjust-mss 1360 R2(config-if)#tunnel source 192.168.20.1 R2(config-if)#tunnel destination 10.10.1.1

Correct Answer: *B*

Community vote distribution

A (67%) B (33%)

⊟ **ChillingAgain** Highly Voted 1 year, 1 month ago

I think we are missing some info in the question to correctly answer this one.

A.
```
R2(config)#interface Tunnel1
R2(config-if)#ip address 172.20.1.2 255.255.255.0
R2(config-if)#ip mtu 1400
R2(config-if)#ip tcp adjust-mss 1360
R2(config-if)#tunnel source 192.168.20.1
R2(config-if)#tunnel destination 192.168.10.1
```

B.
```
R2(config)#interface Tunnel1
R2(config-if#ip address 172.20.1.2 255.255.255.0
R2(config-if)#ip mtu 1400
R2(config-if)#ip tcp adjust-mss 1360
R2(config-if)#tunnel source 10.10.2.2
R2(config-if)#tunnel destination 10.10.1.1
```

C.
```
R2(config)#interface Tunnel1
R2(config-if)#ip address 172.20.1.2 255.255.255.0
R2(config-if)#ip mtu 1500
R2(config-if)#ip tcp adjust-mss 1360
R2(config-if)#tunnel source 10.10.2.2
R2(config-if)#tunnel destination 10.10.1.1
```

D.
```
R2(config)#interface Tunnel1
R2(config-if)#ip address 172.20.1.2 255.255.255.0
R2(config-if)#ip mtu 1500
R2(config-if)#ip tcp adjust-mss 1360
R2(config-if)#tunnel source 192.168.20.1
R2(config-if)#tunnel destination 10.10.1.1
```

upvoted 11 times

⊟ **Mohammad963** Most Recent 4 months ago

I think we cannot answer this Q without exhibit, the only different is with IP address.... please correct me if I'm wrong

upvoted 2 times

⊟ **[Removed]** 5 months ago

wouldn't option A or B be correct?

upvoted 2 times

⊟ **inteldarvid** 5 months, 2 weeks ago

Selected Answer: A

https://community.cisco.com/t5/networking-knowledge-base/how-to-configure-a-gre-tunnel/ta-p/3131970

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: A

sorry option A

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

Selected Answer: B

team for me is B, because, I prefer use Ip class "C" 192.168.1.0/24 for tunnel of gre. I think we are missing some info in the question to correctly answer this one

upvoted 1 times

HungarianDish 7 months, 2 weeks ago

Agree, it's "A" or "B". tunnel source Loopback0 (or the IP) + destination is the IP of the loopback interface on the other end of the tunnel.

upvoted 2 times

DUBC89x 1 year ago

Agreeded, but the answer has to be A or B "IP MTU 1400"

upvoted 1 times

Question #234 *Topic 1*

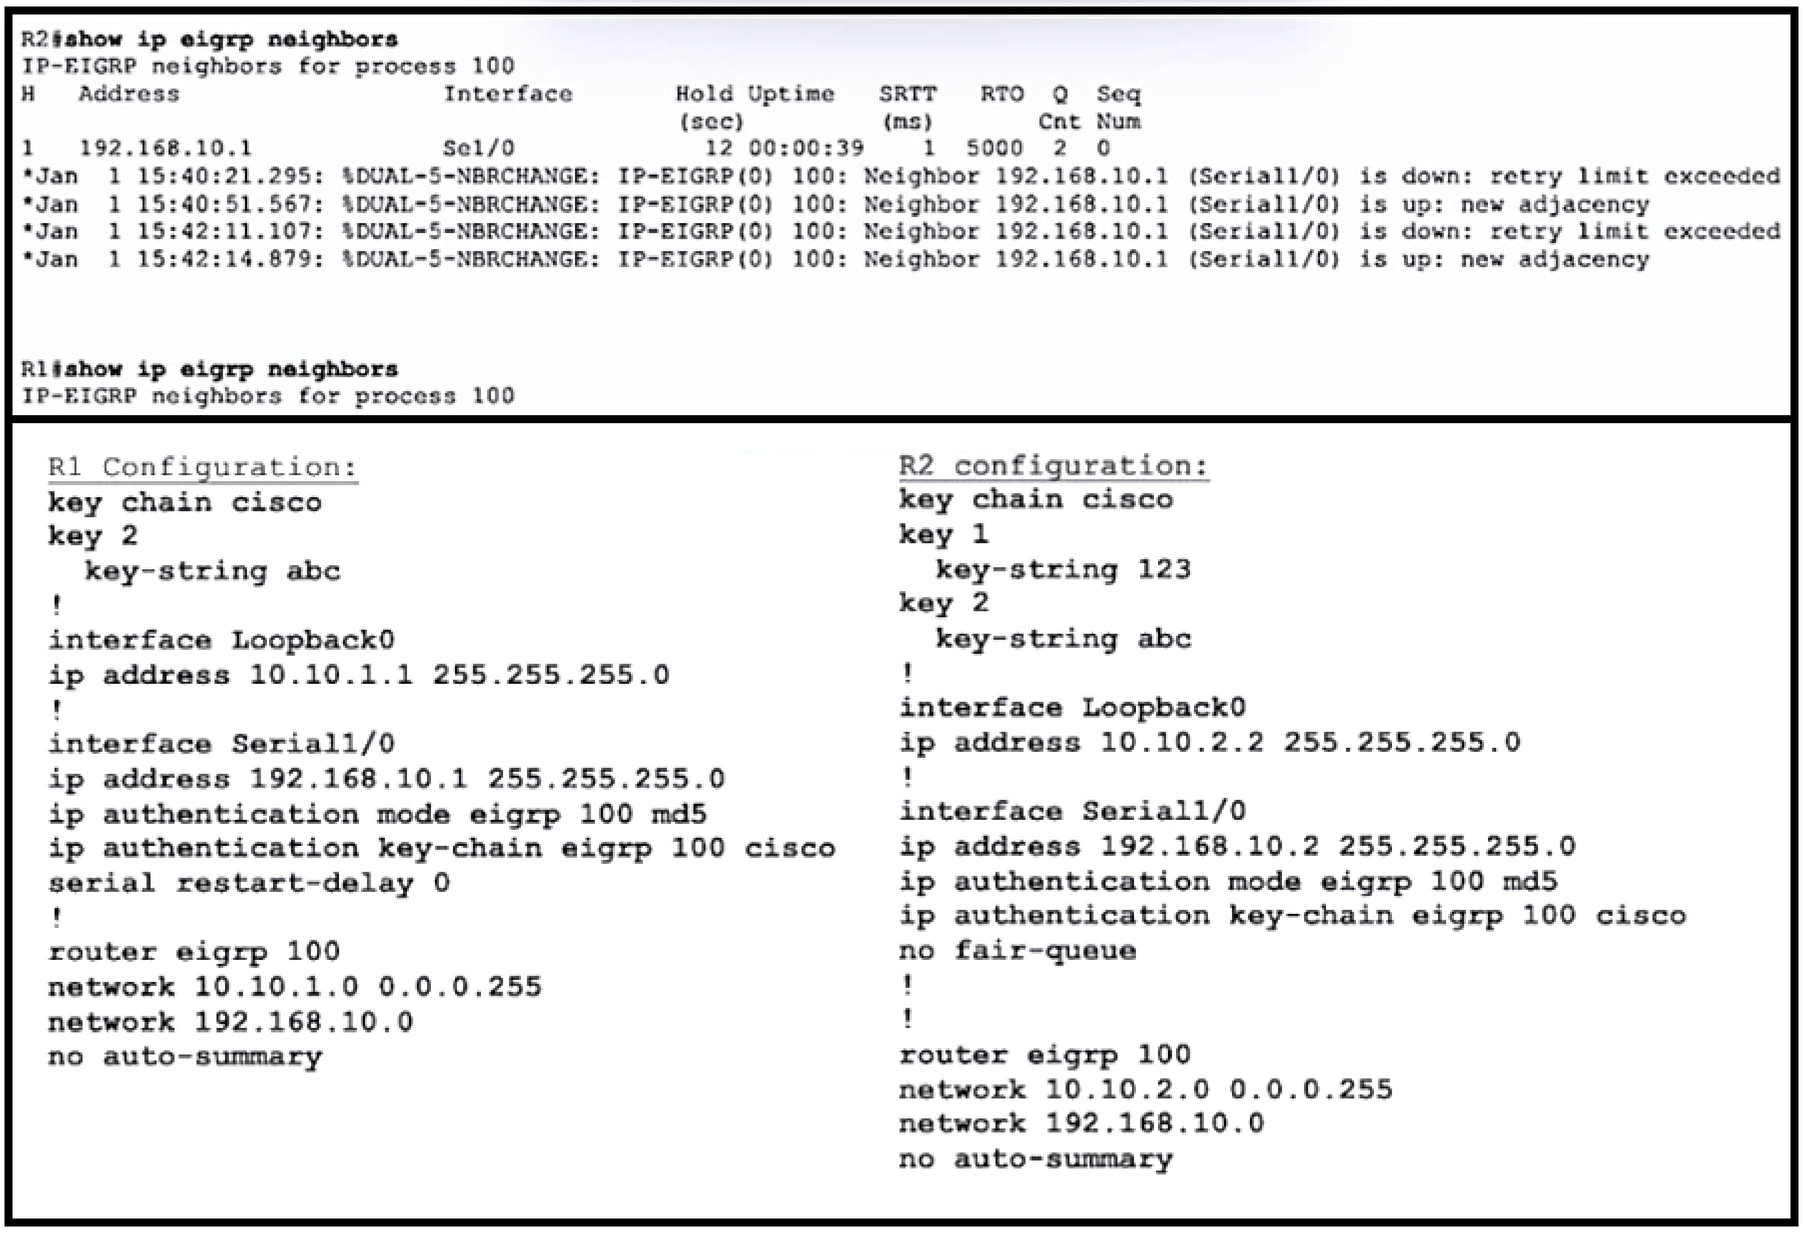

```
R2#show ip eigrp neighbors
IP-EIGRP neighbors for process 100
H   Address                  Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                             (sec)         (ms)       Cnt Num
1   192.168.10.1             Se1/0             12 00:00:39    1   5000  2  0
*Jan  1 15:40:21.295: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 192.168.10.1 (Serial1/0) is down: retry limit exceeded
*Jan  1 15:40:51.567: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 192.168.10.1 (Serial1/0) is up: new adjacency
*Jan  1 15:42:11.107: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 192.168.10.1 (Serial1/0) is down: retry limit exceeded
*Jan  1 15:42:14.879: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 192.168.10.1 (Serial1/0) is up: new adjacency


R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 100
```

R1 Configuration:

```
key chain cisco
key 2
  key-string abc
!
interface Loopback0
ip address 10.10.1.1 255.255.255.0
!
interface Serial1/0
ip address 192.168.10.1 255.255.255.0
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 cisco
serial restart-delay 0
!
router eigrp 100
network 10.10.1.0 0.0.0.255
network 192.168.10.0
no auto-summary
```

R2 configuration:

```
key chain cisco
key 1
  key-string 123
key 2
  key-string abc
!
interface Loopback0
ip address 10.10.2.2 255.255.255.0
!
interface Serial1/0
ip address 192.168.10.2 255.255.255.0
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 cisco
no fair-queue
!
!
router eigrp 100
network 10.10.2.0 0.0.0.255
network 192.168.10.0
no auto-summary
```

Refer to the exhibit. R1 and R2 are configured for EIGRP peering using authentication and the neighbors failed to come up. Which action resolves the issue?

A. Configure a matching lowest key-id on both routers.

B. Configure a matching authentication type on both routers.

C. Configure a matching key-id number on both routers.

D. Configure a matching key-chain name on both routers.

Correct Answer: *A*

Community vote distribution

A (76%) C (24%)

HungarianDish Highly Voted 7 months, 2 weeks ago

Selected Answer: A

For me it's "A". The lowest key ID needs to match, because EIGRP checks against the FIRST valid key. Good sources from you guys:
https://community.cisco.com/t5/routing/eigrp-authentication-problem-need-your-help/td-p/1714446
https://community.cisco.com/t5/switching/key-chain-validation-for-eigrp/td-p/1988487

upvoted 7 times

ZamanR Most Recent 5 days, 20 hours ago

I think A

upvoted 1 times

inteldarvid 5 months, 2 weeks ago

the option correct is A. I test in my lab. Its neecsary put order key chain

upvoted 1 times

HamzaBadar 5 months, 3 weeks ago

Test with GNS3. Answer is A.

upvoted 1 times

Malasxd 7 months, 2 weeks ago

The EIGRP try to use the key-id in the order they were configured. In this exemplo they will never match and there is no way to chance the order EIGRP process the keys.

C seeems more correct for me.

upvoted 1 times

ellen_AA 11 months ago

Selected Answer: A

A is the answer, matching the lowest key-ids on both routers.
https://community.cisco.com/t5/routing/eigrp-authentication-problem-need-your-help/td-p/1714446

upvoted 3 times

JKStinn 11 months, 3 weeks ago

Selected Answer: C

https://community.cisco.com/t5/switching/key-chain-validation-for-eigrp/td-p/1988487

upvoted 1 times

heeeeyajoke 1 year ago

LABBED IT, definitely C

upvoted 1 times

[Removed] 1 year ago

Selected Answer: C

Key numbers need to match

upvoted 1 times

DUBC89x 1 year ago

Selected Answer: C

I reviewed our production environment and there are matching keys.

```
key chain "example"
key 170
key-string "password"
accept-lifetime 10:00:00 Feb 3 2020 infinite
```

upvoted 1 times

Huntkey 1 year, 2 months ago

Selected Answer: A

Although can't find proof on the Internet, I tried and the lowest key ID seems to be a requirement. I would go with A.

upvoted 3 times

jucevabe 1 year, 2 months ago

Selected Answer: C

Answer C

upvoted 1 times

Question #235 *Topic 1*

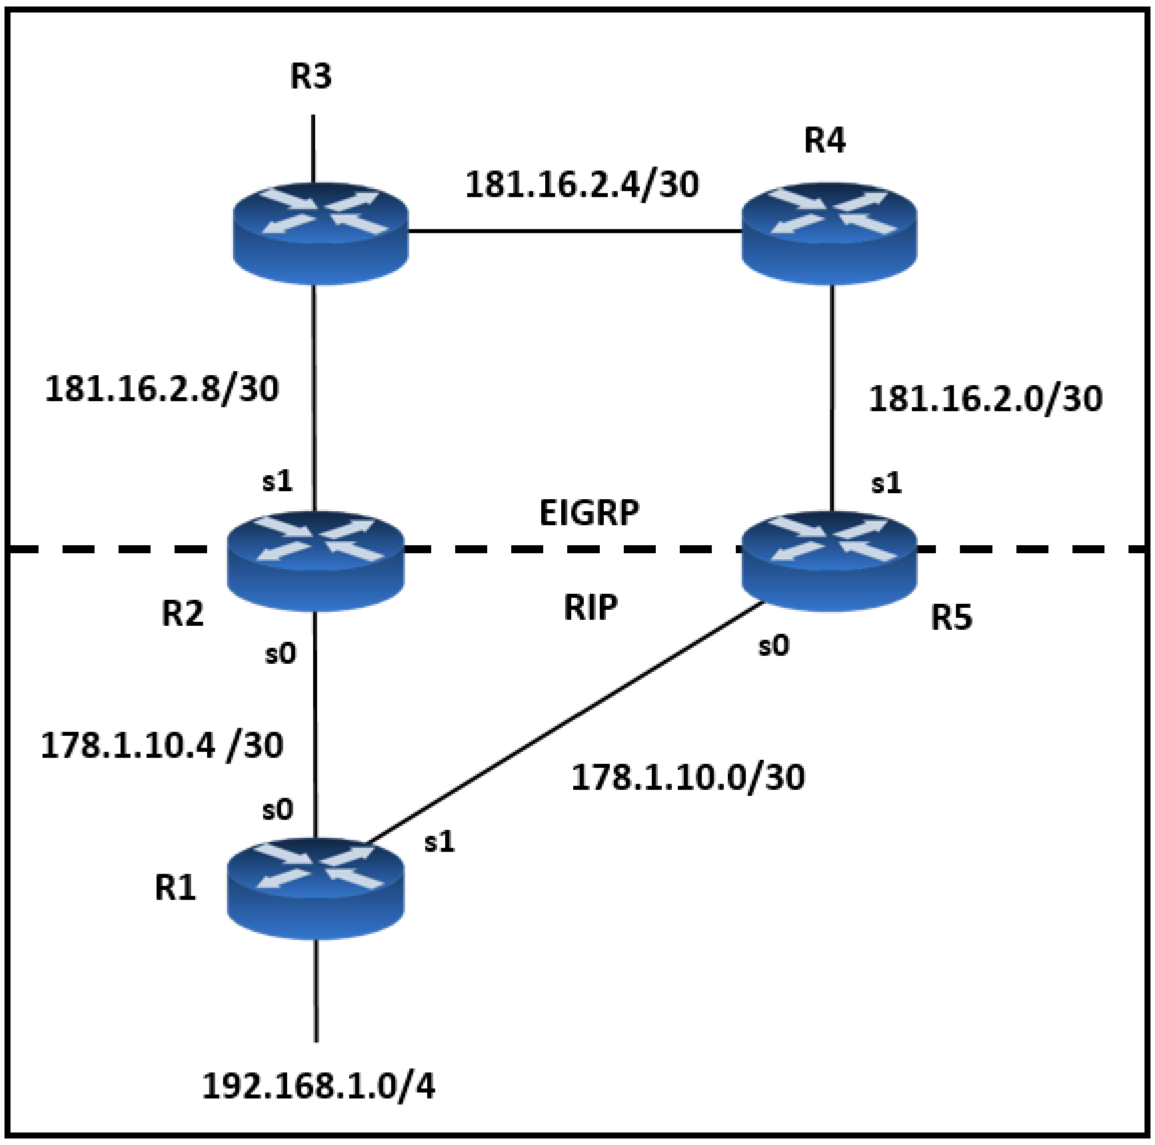

Refer to the exhibit. Mutual redistribution is enabled between RIP and EIGRP on R2 and R5. Which configuration resolves the routing loop for the 192.168.1.0/24 network?

A. R2: router eigrp 10 network 181.16.0.0 redistribute rip metric 1 1 1 1 1 distribute-list 1 in s0 ! router rip network 178.1.0.0 redistribute eigrp 10 metric 2 ! access-list 1 deny 192.168.1.0 access-list 1 permit any R5: router eigrp 10 network 181.16.0.0 redistribute rip metric 1 1 1 1 1 distribute-list 1 in s0 ! router rip network 178.1.0.0 redistribute eigrp 10 metric 2 ! access-list 1 deny 192.168.1.0 access-list 1 permit any

B. R2: router eigrp 10 network 181.16.0.0 redistribute rip metric 1 1 1 1 1 distribute-list 1 in s0 ! router rip network 178.1.0.0 redistribute eigrp 10 metric 2 ! access-list 1 deny 192.168.1.0 access-list 1 permit any R5: router eigrp 10 network 181.16.0.0 redistribute rip metric 1 1 1 1 1 distribute-list 1 in s1 ! router rip network 178.1.0.0 redistribute eigrp 10 metric 2 ! access-list 1 deny 192.168.1.0 access-list 1 permit any

C. R2: router eigrp 10 network 181.16.0.0 redistribute rip metric 1 1 1 1 1 distribute-list 1 in s1 ! router rip network 178.1.0.0 redistribute eigrp 10 metric 2 ! access-list 1 deny 192. 168.1.0 access-list 1 permit any R5: router eigrp 10 network 181.16.0.0 redistribute rip metric 1 1 1 1 1 distribute-list 1 in s0 ! router rip network 178.1.0.0 redistribute eigrp 10 metric 2 ! access-list 1 deny 192.168.1.0 access-list 1 permit any

D. R2: router eigrp 7 network 181.16.0.0 redistribute rip metric 1 1 1 1 1 distribute-list 1 in s1 ! router rip network 178.1.0.0 redistribute eigrp 7 metric 2 ! access-list 1 deny 192.168.1.0 access-list 1 permit any R5: router eigrp 7 network 181.16.0.0 redistribute rip metric 1 1 1 1 1 distribute-list 1 in s1 ! router rip network 178.1.0.0 redistribute eigrp 7 metric 2 ! access-list 1 deny 192.168.1.0 access-list 1 permit any

Correct Answer: *D*

Community vote distribution

D (100%)

[Removed] 4 months, 2 weeks ago

Selected Answer: D

First off, the exhibit is wrong in that a loop is caused in this scenario. EIGRP has two ways of preventing this loop, 1) Split Horizon and 2) EIGRP External routes have an AD of 170.

IGRP would cause a loop because AD is 100 for both internal and external routes.

Life of the route 192.168.1.0/24
1) R1 advertises 192.168.1.0/24 via RIP with AD 120.

2) R2 and R5 learn the route via their links to R1 on RIP with AD 120
3) R2 and R5 redistribute the route into IGRP outbound of interface S1.
4) R3 and R4 learn the route via their links to R2 and R5 respectively and advertise it to each other, R3 to R4, and R4 to R3
5) R2 and R5 learn the route again via their links to R3 and R4, respectively. Note that this route is now learned via IGRP with an AD of 100 which is preferred over RIP AD 120.
6) Loopty doop.

Solution:
Filter the route from being learned via Interface S1 on R2 and R5.

upvoted 1 times

Malasxd 7 months ago

Selected Answer: D

This question with EIGRP does not make any sense. This scenario does not create a looping.
If you replace EIGRP by IGRP the looping is true, just like the exemplo in this link: https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html

With IGRP instead EIGRP, D is right.

upvoted 3 times

[Removed] 4 months, 2 weeks ago

I'm glad I wasn't going crazy, I wrote down notes on how the route would be installed in the RIB. R2 and R5 will install the route as RIP AD 120, which will be preferred over EIGRP external AD 170, R3 and R4 will not have a cause for Loop as they will learn it through R2 and R5 respectively with AD 170.

No loop here.

upvoted 2 times

HungarianDish 6 months, 2 weeks ago

Good point! The example from the above article is for IGRP - RIP redistribution. IGRP has AD of 100 for both internal and external routes. So, R2 and R5 are going to prefer the IGRP path for 192.168.1.0/24, and not RIP with AD 120. EIGRP with external AD of 170 won't have this issue. Btw, the article has an example for EIGRP "Example 2". I think that their "Example 2" is not entirely correct.)

upvoted 2 times

HungarianDish 6 months, 2 weeks ago

Some questions have a really poor quality. Based on experience, it is not better on the real exam either. :(

upvoted 2 times

HungarianDish 7 months, 2 weeks ago

Selected Answer: D

Example is taken from the cisco document linked by BECAUSE:
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html
"[R2 and R5] are told that they must not learn network 192.168.1.0/24 through the EIGRP updates they receive on their serial 1 interface. Therefore, the only knowledge these routers have for network 192.168.1.0/24 is through RIP from R1."

upvoted 2 times

HungarianDish 7 months, 2 weeks ago

```
R2
router igrp 7
network 172.16.0.181
redistribute rip metric 1 1 1 1 1
distribute-list 1 in s1

router rip
network 172.16.0.0
redistribute igrp 7 metric 2

access-list 1 deny 192.168.1.0
access-list 1 permit any

R5
router igrp 7
network 172.16.0.181
redistribute rip metric 1 1 1 1 1
distribute-list 1 in s1

router rip
network 172.16.0.0
redistribute igrp 7 metric 2

access-list 1 deny 192.168.1.0
access-list 1 permit any
```

upvoted 2 times

DUBC89x 1 year ago

I agree D
Both R2 and R5 are redistributing the rip route for 192.168.1.0/24. What you want to do is block that route that is being received from R3/R4 and being redistributed back into R2 and R5.

upvoted 1 times

BECAUSE 1 year, 2 months ago

Selected Answer: D

Given answer is correct.
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html

upvoted 2 times

Question #236

Topic 1

Which method provides failure detection in BFD?

A. long duration, low overhead

B. short duration, low overhead

C. long duration, high overhead

D. short duration, high overhead

Correct Answer: *B*

Community vote distribution

B (100%)

HungarianDish 7 months, 2 weeks ago

Selected Answer: B

"BFD Operation
BFD provides a low-overhead, short-duration method of detecting failures in the forwarding path between two adjacent routers, including the interfaces, data links, and forwarding planes. BFD is a detection protocol that you enable at the interface and routing protocol levels."
https://www.cisco.com/en/US/docs/ios/12_4t/ip_route/configuration/guide/t_bfd.html

upvoted 4 times

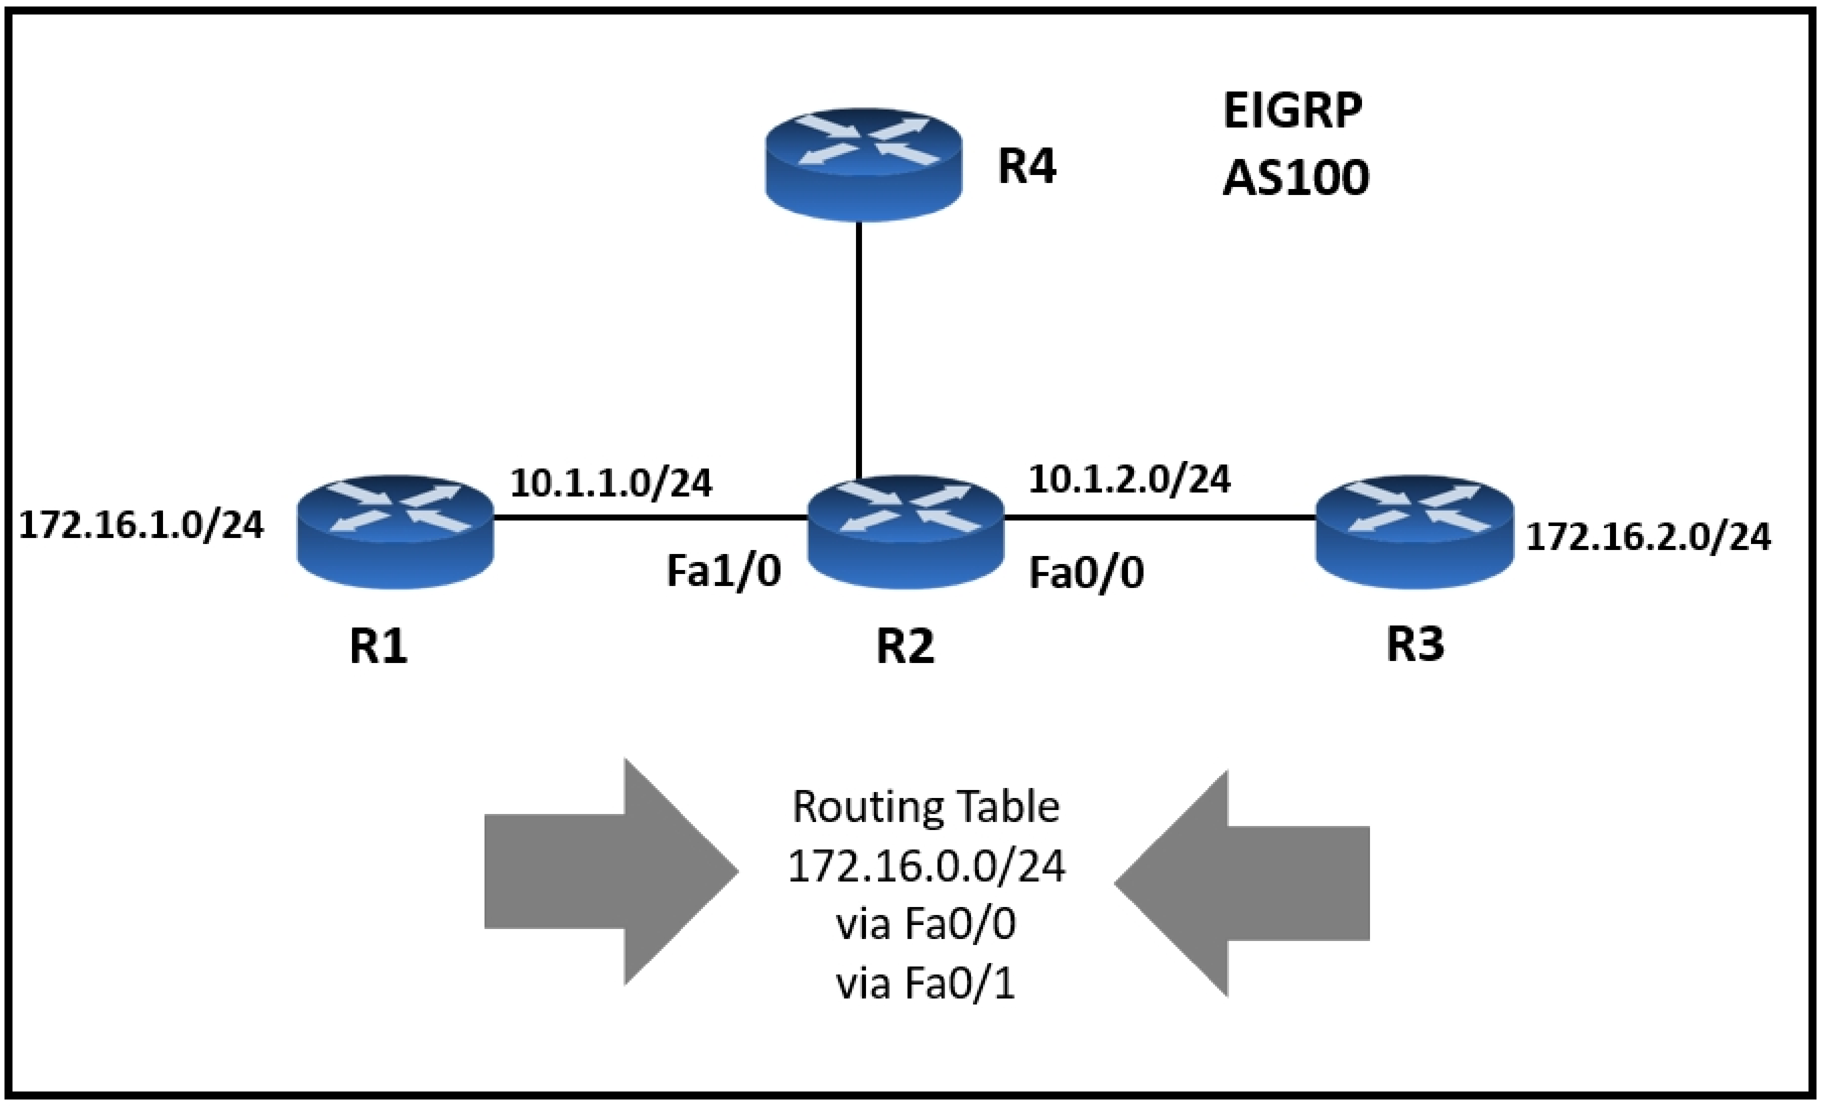

Refer to the exhibit. R4 is experiencing packet drop when trying to reach 172.16.2.7 behind R2. Which action resolves the issue?

A. Insert a /24 floating static route on R2 toward R3 with metric 254.

B. Disable auto summarization on R2.

C. Enable auto summarization on all three routers R1, R2, and R3.

D. Insert a /16 floating static route on R2 toward R3 with metric 254.

Correct Answer: *B*

HungarianDish 7 months, 2 weeks ago
With auto-summary enabled, subnets will be advertised as classful networks. This causes problems with discontiguous networks. R2 will think it has two equal paths (via R1 and R3) to reach 172.16.0.0/16.
https://networklessons.com/eigrp/eigrp-auto-summary
upvoted 2 times

heeeeyajoke 1 year ago
i believe its a /16 thats currently in the routing table, even with a 172.16.0.0 /24 route, the router is still not aware of the existence of the interesting route, so creating the route is still valid
upvoted 1 times

heeeeyajoke 1 year ago
This is supposed to be A, the /24 prefix is being sent to R2 by its neighbours, the only way it will route properly to the desired /24 prefix is to create a route with a longer prefix. This takes precedence over the current summarized route in the routing table
upvoted 2 times

Pietjeplukgeluk 4 days, 21 hours ago
Indeed, R2 will receive a summarized route, so B will not be able to undo this.
upvoted 1 times

Pietjeplukgeluk 4 days, 20 hours ago
I checked https://networklessons.com/cisco/ccie-routing-switching/eigrp-auto-summary and for sure the answer is A (!) here. Disable auto summary only works on the routers that actually advertise the routes initialy.
upvoted 1 times

Question #238

Topic 1

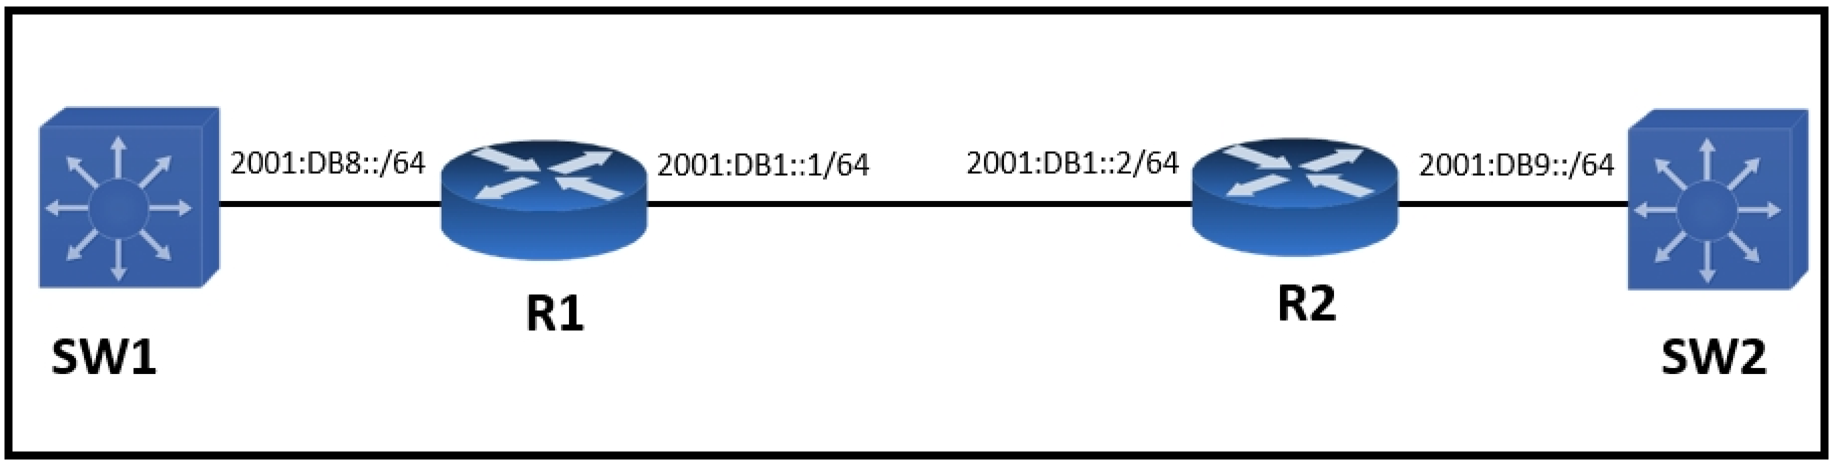

Refer to the exhibit. An engineer must advertise routes into IPv6 MP-BGP and failed. Which configuration resolves the issue on R1?

A. router bgp 64900 no bgp default ipv4-unicast address-family ipv6 unicast redistribute ospf network 2001:DB9::/64

B. router bgp 64900 no bgp default ipv4-unicast address-family ipv6 multicast neighbor 2001:DB8:7000::2 translate-update ipv6 multicast

C. router bgp 65000 no bgp default ipv4-unicast address-family ipv6 unicast network 2001:DB8::/64

D. router bgp 65000 no bgp default ipv4-unicast address-family ipv6 multicast network 2001:DB8::/64

Correct Answer: *C*

Community vote distribution

C (100%)

HungarianDish 7 months, 2 weeks ago

Selected Answer: C

"The command is an enabler for Multi protocol BGP mode where multiple address families can be negotiated during the BGP session setup..."
"The need for this command "no bgp default ipv4-unicast" may have been removed in recent IOS images by reverting the default BGP behaviuor to be Multi protocol."
https://community.cisco.com/t5/routing/no-bgp-default-ipv4-unicast/td-p/2913083
https://community.cisco.com/t5/mpls/quot-no-bgp-default-ipv4-unicast-quot-command/td-p/1212139

upvoted 1 times

heeeeyajoke 1 year ago

Answer is correct

upvoted 1 times

Question #239

Topic 1

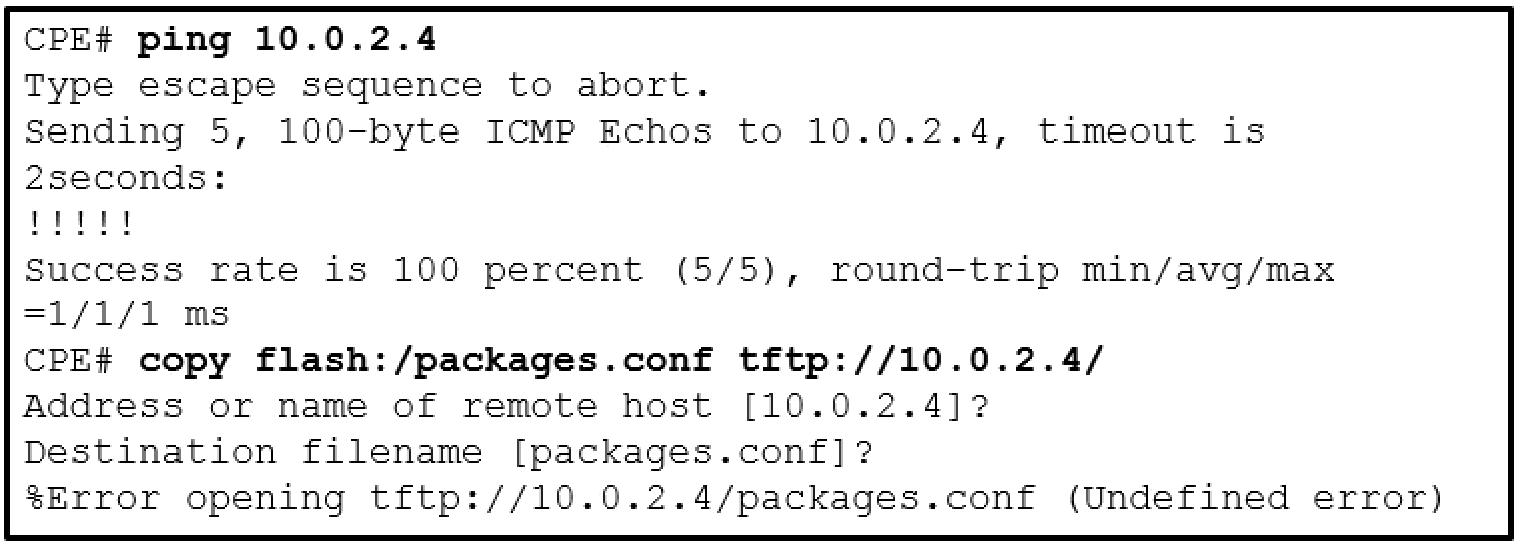

```
CPE# ping 10.0.2.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.0.2.4, timeout is
2seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max
=1/1/1 ms
CPE# copy flash:/packages.conf tftp://10.0.2.4/
Address or name of remote host [10.0.2.4]?
Destination filename [packages.conf]?
%Error opening tftp://10.0.2.4/packages.conf (Undefined error)
```

Refer to the exhibit. The administrator is trying to overwrite an existing file on the TFTP server that was previously uploaded by another router. However, the attempt to update the file fails.

Which action resolves this issue?

A. Make the TFTP folder writable by all on the TFTP server.

B. Make the package.conf file writable by all on the TFTP server.

C. Make the package.conf file executable by all on the TFTP server.

D. Make sure to run the TFTP service on the TFTP server.

Correct Answer: *B*

Community vote distribution

B (57%) A (43%)

ZamanR 5 days, 7 hours ago

B is correct answer

upvoted 1 times

guy276465281819372 5 months ago

Selected Answer: A

I believe the question is incorrect and misleading,
It depends on the file system that the tftp server run on.
I would go for A anyway.

upvoted 3 times

Pietjeplukgeluk 4 days, 19 hours ago

When Linux is used changing the folder permission only does not make a difference on files within (except when applying them recursively of course), also on windows you can actually disable inheritance of rights. Anyway, the only solution that always works is allowing the file to be written. So in my understanding it is B. To solve future use cases i would also change the folder rights, but that is actually not the question here.

upvoted 1 times

[Removed] 5 months ago

Selected Answer: B

The key to the question is the phrase "to overwrite an existing file on the TFTP server". We can only assume that the file is the same name, and if the TFTP server does not allow the file that already exists to be rewriteable then an error ocurrs.

upvoted 4 times

Me_3e 4 months, 2 weeks ago

agree because "that was previously uploaded by another router" seem user can writable in the folder but .conf is not sure.

upvoted 1 times

Question #240 *Topic 1*

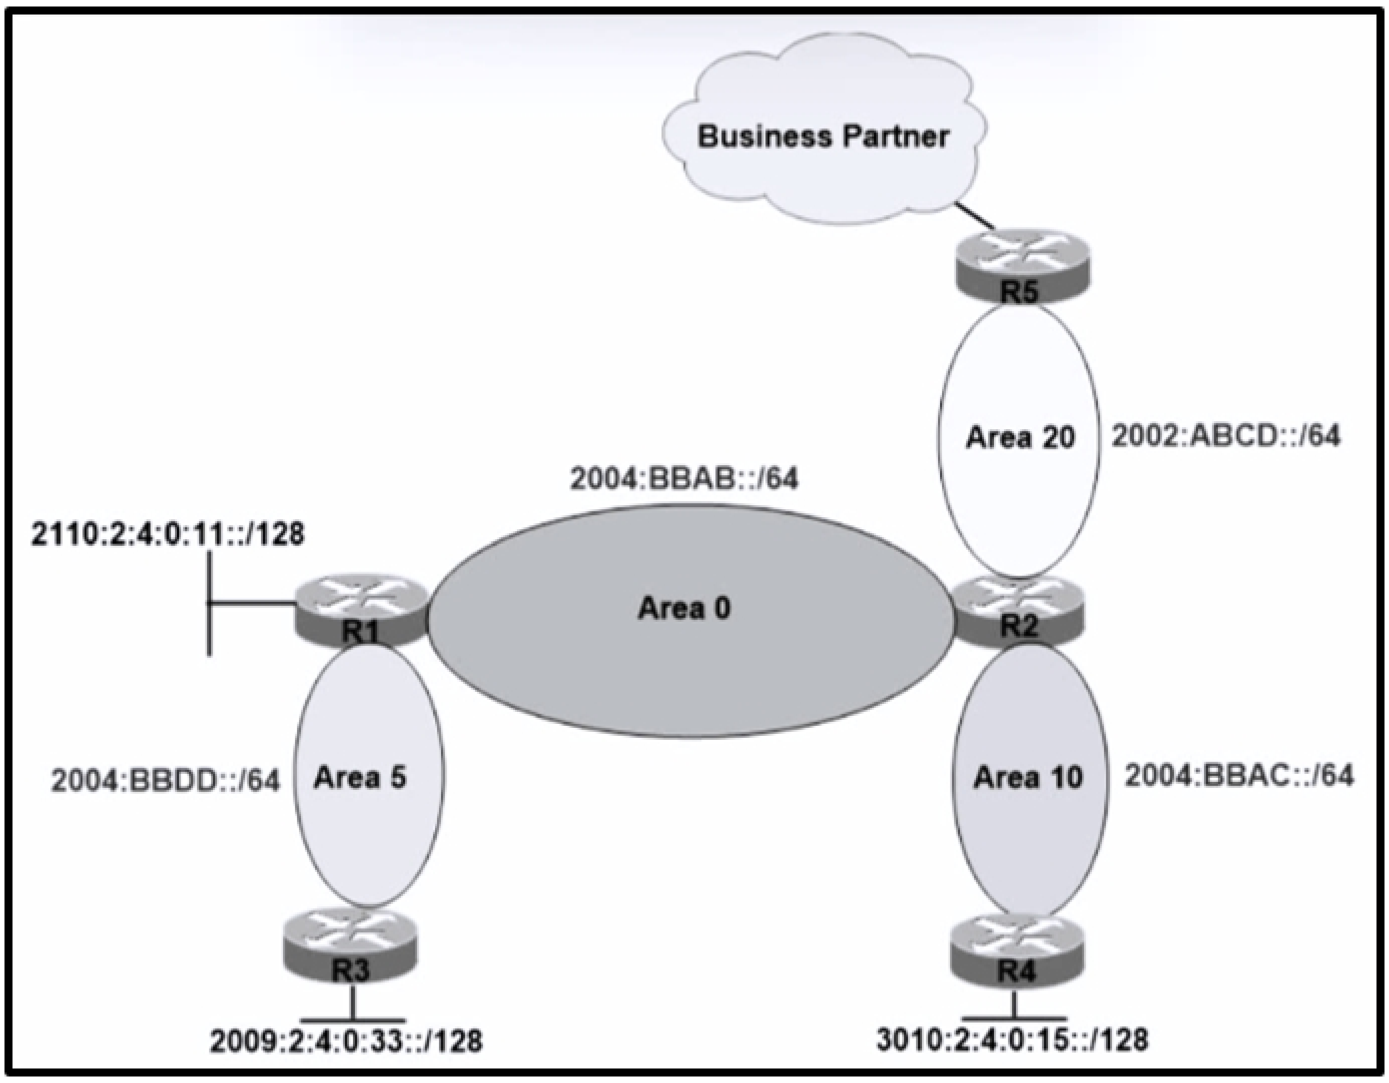

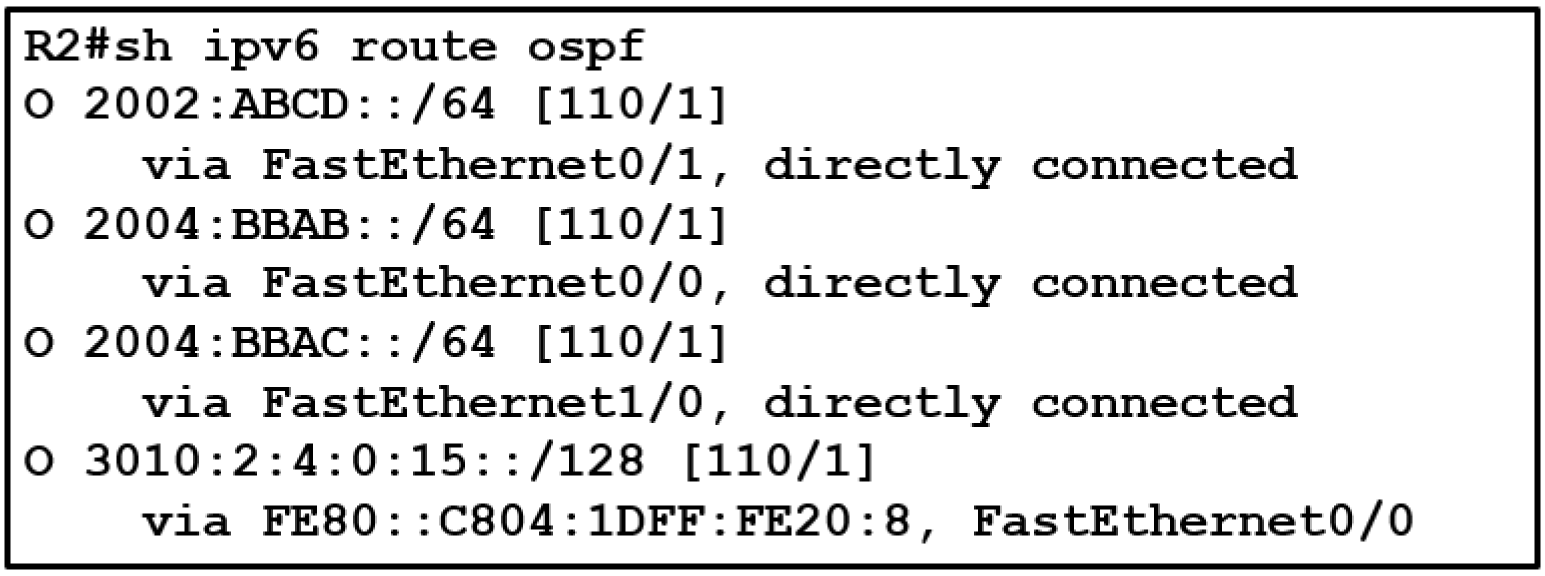

```
R2#sh ipv6 route ospf
O 2002:ABCD::/64 [110/1]
    via FastEthernet0/1, directly connected
O 2004:BBAB::/64 [110/1]
    via FastEthernet0/0, directly connected
O 2004:BBAC::/64 [110/1]
    via FastEthernet1/0, directly connected
O 3010:2:4:0:15::/128 [110/1]
    via FE80::C804:1DFF:FE20:8, FastEthernet0/0
```

Refer to the exhibit. A network engineer applied a filter for ISA traffic on OSPFv3 inter area routes on the area 5 ABR to protect advertising the internal routes of area 5 to the business partner network. All other areas should receive the area 5 internal routes. After the respective route filtering configuration is applied on the
ABR, area 5 routes are not visible on any of the areas. How must the filter list be applied on the ABR to resolve this issue?

A. in the "in" direction for area 5 on router R1

B. in the "in" direction for area 20 on router R2

C. in the "out" direction for area 20 on router R2

D. in the "out" direction for area 5 on router R1

Correct Answer: *D*

Community vote distribution

C (52%) B (48%)

tamangao Highly Voted 1 year, 1 month ago

B is the right answer, lab it.

upvoted 8 times

bryaberson 2 months, 4 weeks ago

R5 is not an ABR. Question states the conf must be applied to the ABR which is R2.
Answer is C

upvoted 1 times

⊟ **asans** Most Recent 2 days, 2 hours ago

For those that say C, when you filter OUT of Area 20 on R2, to which areas are you preventing the updates to enter? That would be Areas 0, 10 which is not what's required here. The thing is on which Router are you filtering. OUT direction would work if the filtering was being done on R5. The correct answer is B, you are filtering INto Area 20 so that updates from Area 0, 5 and Area 10 are not allowed in 20.

upvoted 1 times

⊟ **MJM1973** 3 weeks, 1 day ago

B is correct answer

Logic is to pay attention to the option in which area it has asked to apply a filter. If the option was for Area 0 then it should be in the OUT direction i.e. routes going out of Area 0 into another area.

The question has asked to apply a filter in Area 20 which can satisfy condition - inter area routes on the area 5 ABR should be protected from getting into the business partner network, all other areas (area 00 and Area 10) should receive the area 5 internal routes
R5 is the only router in Area 20 and Area 20 needs to be blocked from receiving routes hence it should be "IN" Direction i.e routes should be blocked entering area 20

https://community.cisco.com/t5/networking-knowledge-base/using-area-lt-gt-filter-llist-command-in-ospf/ta-p/3118832

upvoted 1 times

⊟ **Mishranihal737** 2 months, 3 weeks ago

Selected Answer: C

C is the correct answer I don't know how tamangao labbed it.

upvoted 1 times

⊟ **Fenix7** 3 months, 1 week ago

"in" direction to R2 will block area 5 propagating to area 10. So, it's C, "out" R2

upvoted 1 times

⊟ **fizzer** 3 months, 1 week ago

Right answer is B
The question is a bit misleading by saying every area must learn about the route
I believe R2 belongs to us but R5 belongs to the business partner, Area 20 is the area between us and the business partner that we both use to share routes and this is where we control what routes we share with them, by filtering what goes into area 20 on R2

A is just silly trying to filter a route that originates for Area 5 from coming into Area 5

D is the exact thing the Engineer has done which is causing the problem - prevent R1 from advertising the route from area 5 into area 0 which means R2 does not know about it in area 0 to advertise it further to area 10 where it is also needed internally

C is telling R2 not advertise the route from area 20 to any other areas (configured on R2), however, area 20 would need to learn about the route from area 0 first before it can decide whether or not to advertise it out to another area. Best place to apply option C would be on R5 but as mentioned above, I do not believe it is under our jurisdiction

upvoted 1 times

⊟ **Colmenarez** 3 months, 1 week ago

Selected Answer: B

B is correct

upvoted 1 times

⊟ **HarwinderSekhon** 4 months, 1 week ago

Selected Answer: C

I am going with C, Why- We need to have routes in area 20 but block it from going into bussiness partner network. If you filter inbound, you block it from being injected to area20 but question says you need to have it in all areas.

upvoted 1 times

⊟ **guy276465281819372** 5 months ago

Selected Answer: C

c is the answer

upvoted 1 times

⊟ **Juraj22** 5 months, 1 week ago

Yes, But this is not good question again....Area 20 should recieve Area 5 routes..Area 20 is not Buisniss partner network...You should filter on ASBR(R5), or second concept of meaning is OUT direction to Area 20(but not right in prefix meaning)...for me unclear question

upvoted 3 times

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: B

team sorry, my answer before is wrong the answer corret is B. is in for area 20

upvoted 1 times

⊟ **MicMillon** 5 months, 1 week ago

Selected Answer: C

Out direction, on R2

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

team the option correct is "C", because, if block in router area 20 block another area 10. Only block area busisnes.Option C

upvoted 1 times

Juraj22 5 months, 4 weeks ago

Selected Answer: C

100% C

upvoted 2 times

wigola 6 months, 1 week ago

R2(config-router)# area 20 filter-list prefix OSPF-FILTER ?
in Filter networks sent to this area
out Filter networks sent from this area
...so, B is the correct one

upvoted 2 times

slcc99 6 months, 1 week ago

Answer is correct.
"All other areas should receive the area 5 internal routes.", but "After the respective route filtering configuration is applied on the ABR, area 5 routes are not visible on any of the areas.", which suggests that route filtering in the "out" direction is configured on R1.You need to fix the "out" direction route filtering configured on R1.

upvoted 1 times

NetworkVal 6 months, 2 weeks ago

Answer is "B"

upvoted 1 times

Question #241

Topic 1

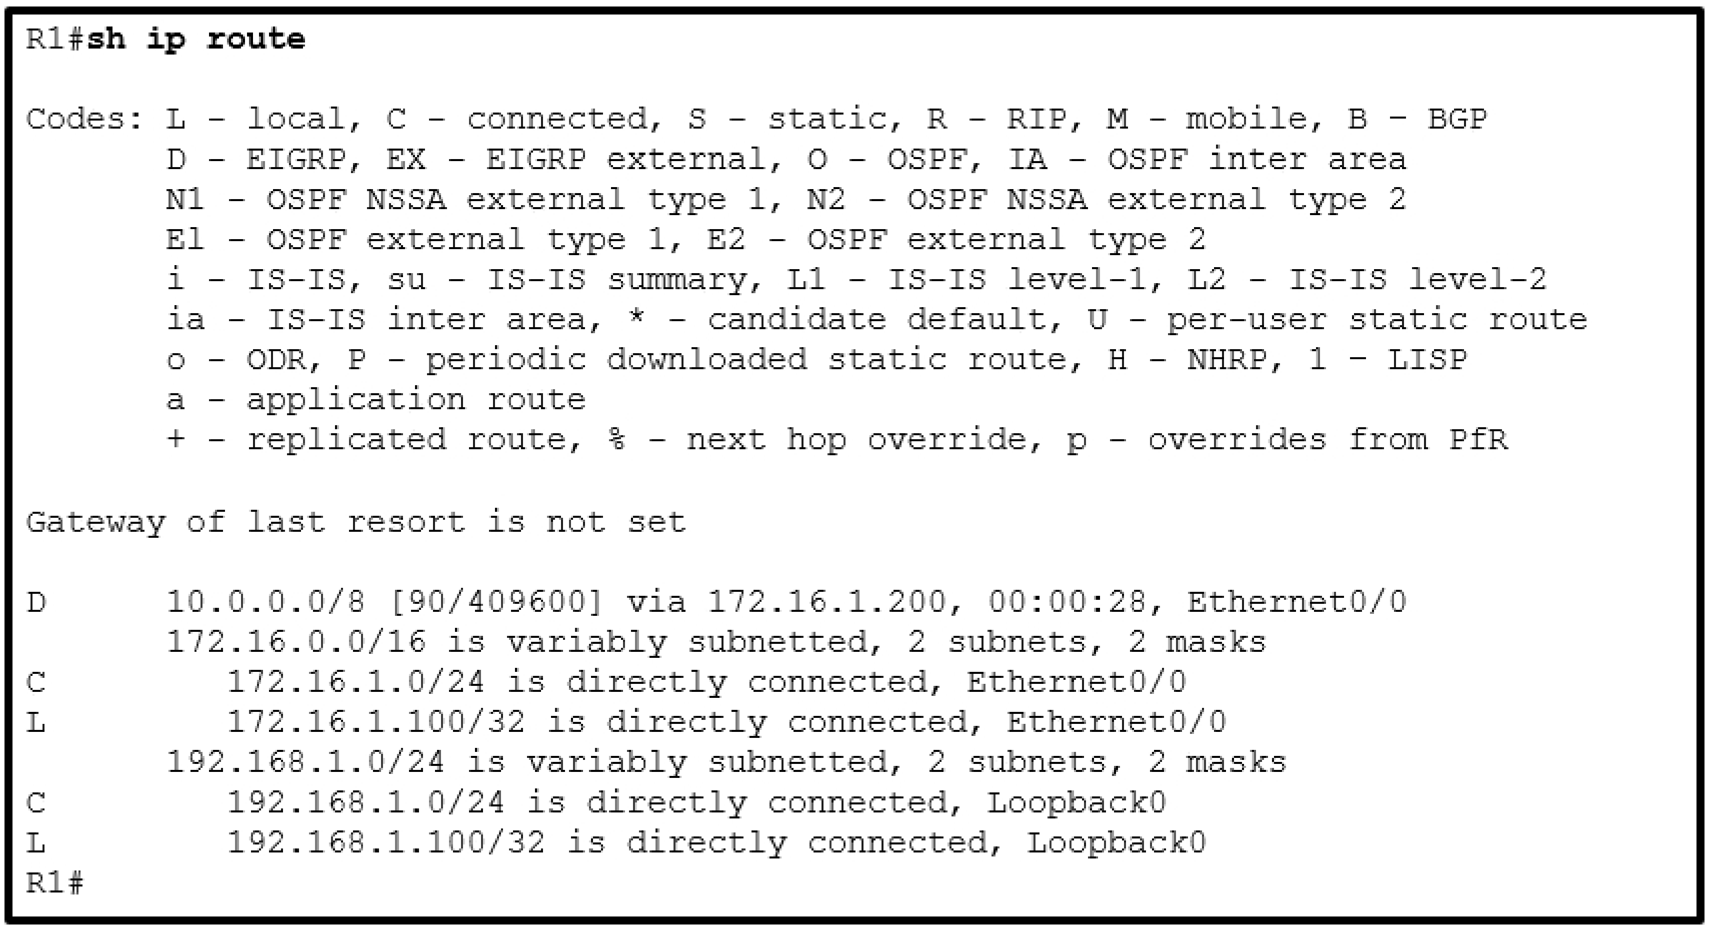

```
R1#sh ip route

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       El - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, 1 - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR

Gateway of last resort is not set

D      10.0.0.0/8 [90/409600] via 172.16.1.200, 00:00:28, Ethernet0/0
       172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C         172.16.1.0/24 is directly connected, Ethernet0/0
L         172.16.1.100/32 is directly connected, Ethernet0/0
       192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C         192.168.1.0/24 is directly connected, Loopback0
L         192.168.1.100/32 is directly connected, Loopback0
R1#
```

Refer to the exhibit. The R2 loopback interface is advertised with RIP and EIGRP using default values. Which configuration changes make R1 reach the R2 loopback using RIP?

A. R1(config)#router rip R1(config-router)#distance 90

B. R1(config)#router eigrp 1 R1(config-router)#distance eigrp 130 120

C. R1(config)#router rip R1(config-router)#distance 100

D. R1(config)#router eigrp 1 R1(config-router)#distance eigrp 120 120

Correct Answer: *B*

Community vote distribution

B (100%)

ChillingAgain Highly Voted 1 year, 1 month ago

Selected Answer: B

"distance eigrp 130 120" set the internal EiGRP routes to 130 and external EIGRP routes to 120. As of the loopback address is advertised in EIGRP as internal route it has an AD of 130. So the RIP route with an AD of 120 is preferred now.

upvoted 5 times

inteldarvid Most Recent 5 months, 1 week ago

Selected Answer: B

yes, teh option correct is B

upvoted 1 times

pepgua 6 months, 1 week ago

Selected Answer: B

A. R1(config)# router rip
R1(config-router)# distance 90
B. R1(config)# router eigrp 1
R1(config-router)# distance eigrp 130 120
C. R1(config)# router rip
R1(config-router)# distance 100
D. R1(config)# router eigrp 1
R1(config-router)# distance eigrp 120 120

upvoted 1 times

chris7890 1 year, 2 months ago

Rip has the administrative distance of 120 to make a clear decision we use a higher / worse AD of 130

upvoted 1 times

Huntkey 1 year, 2 months ago

B is fine but why not D?

upvoted 1 times

mrnipsnips 1 year, 1 month ago

D will make EIGRP and RIP equal

upvoted 2 times

Question #242 Topic 1

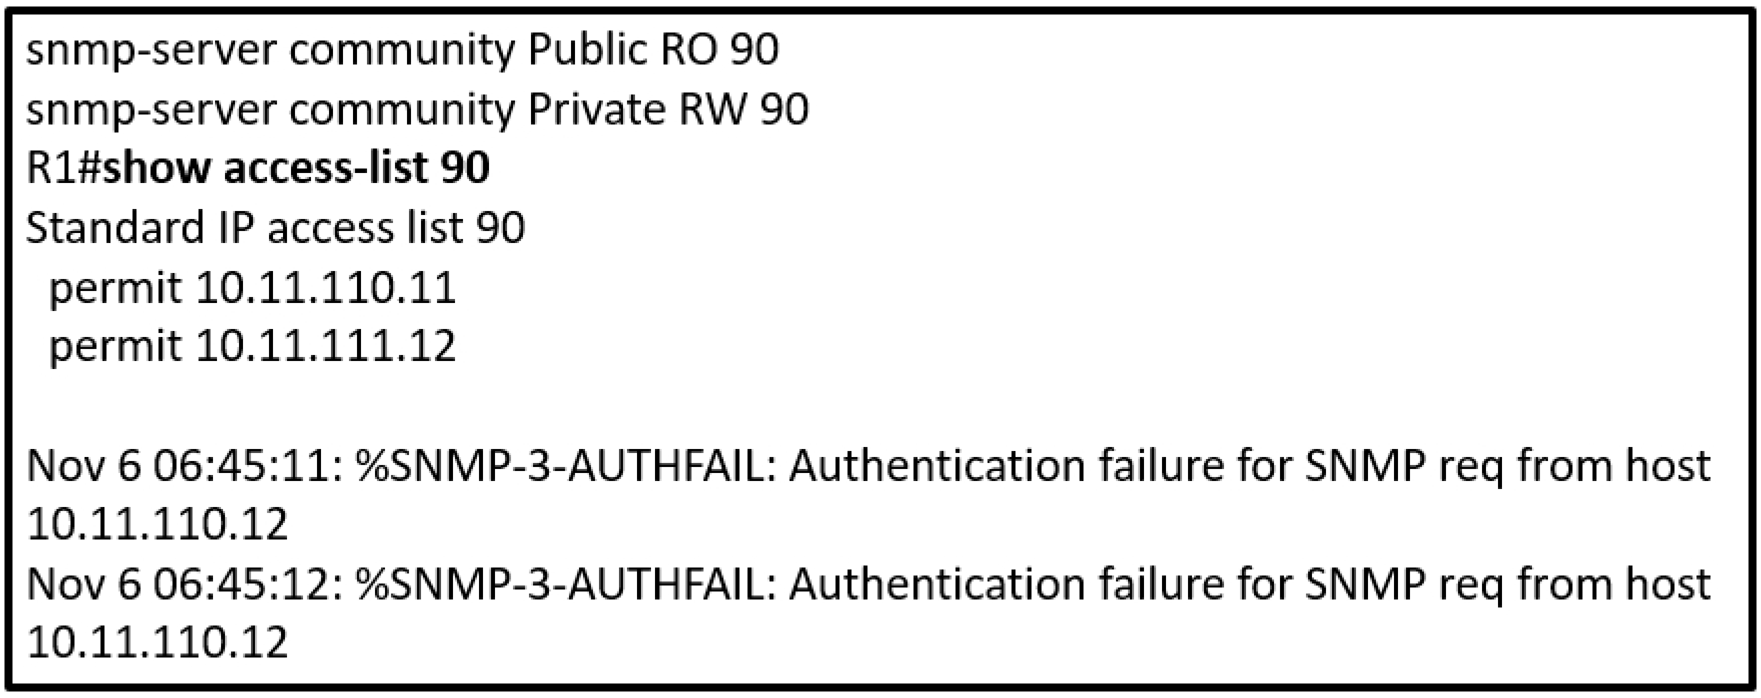

```
snmp-server community Public RO 90
snmp-server community Private RW 90
R1#show access-list 90
Standard IP access list 90
  permit 10.11.110.11
  permit 10.11.111.12

Nov 6 06:45:11: %SNMP-3-AUTHFAIL: Authentication failure for SNMP req from host
10.11.110.12
Nov 6 06:45:12: %SNMP-3-AUTHFAIL: Authentication failure for SNMP req from host
10.11.110.12
```

Refer to the exhibit. A network administrator notices these console messages from host 10.11.110.12 originating from interface E1/0. The administrator considers this an unauthorized attempt to access SNMP on R1. Which action prevents the attempts to reach R1 E1/0?

A. Configure IOS control plane protection using ACL 90 on interface E1/0.

B. Create an inbound ACL on interface E1/0 to deny SNMP from host 10.11.110.12.

C. Add a permit statement including the host 10.11.110.12 into ACL 90.

D. Configure IOS management plane protection using ACL 90 on interface E1/0.

Correct Answer: *B*

Community vote distribution

D (83%) B (17%)

ZamanR 5 days, 8 hours ago

D is correct

upvoted 1 times

Fenix7 3 months, 1 week ago

snmp-server community Public RO 90
snmp-server community Private W 90
R1#show access-list 90
Standard IP access list 90
permit 10.11.110.11
permit 10.11.111.12

Console messages are from 10.11.110.12

See the difference between the permit IP statement and host IP?

B is correct.

upvoted 2 times

[Removed] 4 months, 4 weeks ago

Selected Answer: D

Lets think through this.
A) is wrong because SNMP functions in the management not the control plane.
B) this sounds correct, but if you think about it, it may cause unintended traffic denies. If we create a new ACL to deny the host, the answer does not specify other parameters, and we could assume that a permit any at the end will be configured as well.
C) is wrong, we are trying to block the host.
D) seems to be the best answer. If we use the same ACL 90, we are inherently deny any other hosts that do not require access to R1's management plane, and only permit the ones defined in the ACL.

D is the best answer
B works, but not entirely the best answer.

upvoted 4 times

guy276465281819372 5 months ago

Selected Answer: D

The question does not specify if the new ACL (answer B) will allow other hosts to access the router through E1/0. I believe the best answer would be D as it uses the existing ACL which block access from the suspected attacker to access R1.

upvoted 1 times

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: B

yes, correct option B. Easy question

upvoted 1 times

Question #243 Topic 1

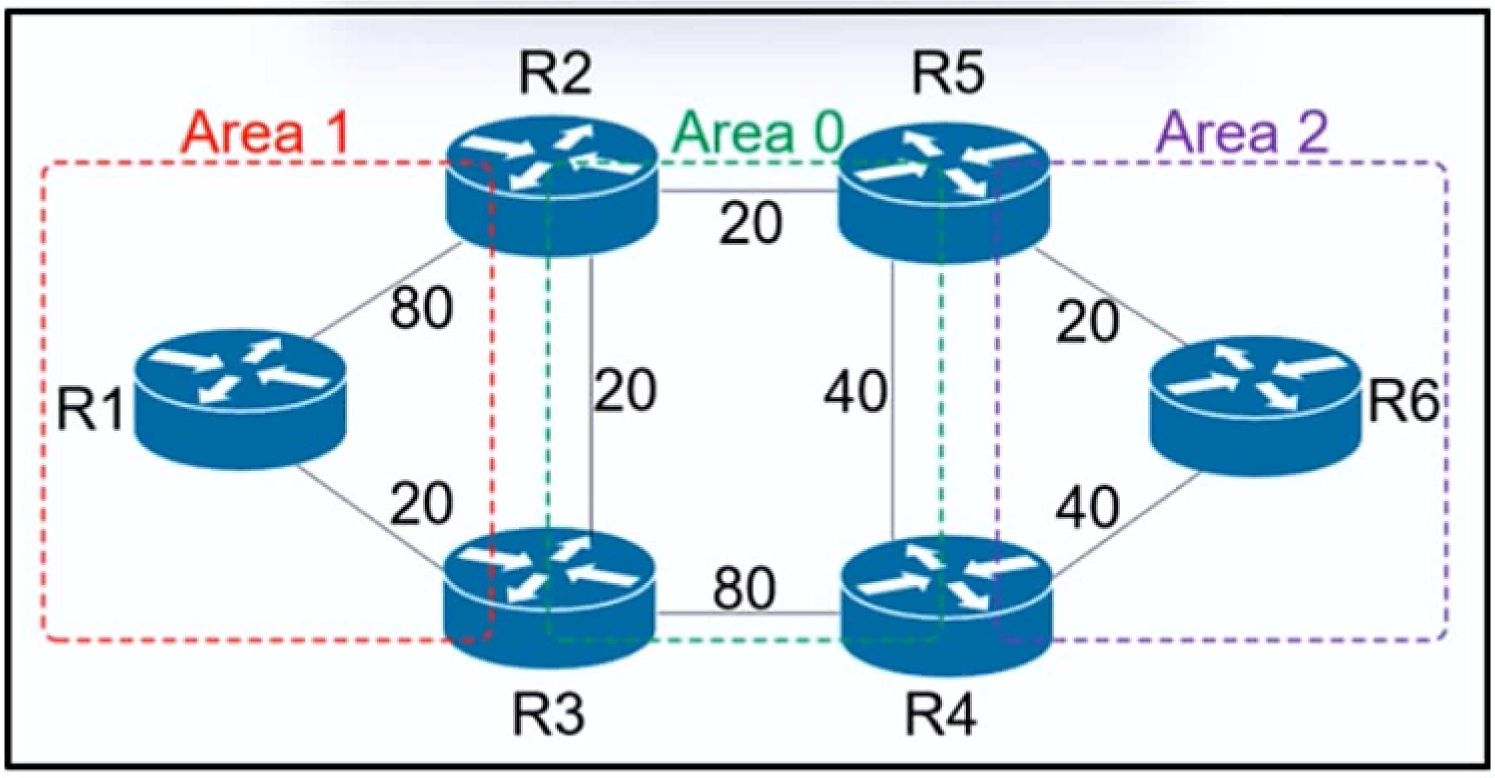

Refer to the exhibit. R6 should reach R1 via R5>R2>R1. Which action resolves the issue?

A. Decrease the cost to 2 between R6-R5-R2.

B. Increase the cost to 61 between R2-R3-R1.

C. Increase the cost to 61 between R2 and R3.

D. Decrease the cost to 41 between R2 and R1.

Correct Answer: *C*

Community vote distribution

C (86%) 14%

[Removed] 4 months, 4 weeks ago

Selected Answer: C

what a stupid ass question, instead of testing your knowledge, the test you on whether or not you can catch the fucking typo or trick word in the answers.
R2 to R3 cost is 20, we can't "decrease" the value to 41, we can increase it. just because of that answer D is wrong.
Answer C has the correct wording.

fuck cisco for letting this shit questions continue.

upvoted 2 times

pepgua 6 months, 1 week ago

Selected Answer: C

Increase cost BY 61. Keyword BY, not TO 61. When you increase BY 61, total becomes 81 which is what you want to achieve.

upvoted 1 times

bucket12678 6 months, 3 weeks ago

I detest how these are worded sometimes. Technically speaking, if you increase the cost TO 61, then the cost of the link = 61 (in which case, it doesn't use the R1-R2 link). However, if you increase the cost BY 61, then the cost of the link = 81 (20+61). This is just splitting hairs over semantics, but the question is worded incorrectly.

upvoted 2 times

HungarianDish 7 months, 2 weeks ago

Selected Answer: C

Agree with Demir11's calculation.

upvoted 3 times

[Removed] 8 months, 2 weeks ago

Basically due to the lowest cost the path it is taking before the change is R6-R5-R2-R3-R1
Increasing the cost by 61 makes the total cost 81>80 so it will prefer R1

upvoted 2 times

6dd4aa0 8 months, 2 weeks ago

Selected Answer: D

See the answer provided by sol_ls95. I agreed with the user.

upvoted 1 times

pulsetion 8 months ago

What sol_ls95 said is incorrect. This would make 6-5-2-1= 81 and 6-5-2-3-1=80.

upvoted 2 times

sol_ls95 11 months, 2 weeks ago

a)6-5-2-1=82 6-5-2-3-1=42
b)6-5-2-1=120 6-5-2-3-1=111
c)6-5-2-1=120 6-5-2-3-1=141
d)6-5-2-1=81 6-5-2-3-1=89

upvoted 1 times

pyrokar 7 months, 1 week ago

The calculation for d is wrong or probably a typo, 6-5-2-3-1=80.
Thus the solution is C.

upvoted 1 times

babs 12 months ago

can someone explain this?

upvoted 1 times

ellen_AA 11 months, 2 weeks ago

On each suggestion, try replace the cost with the suggested one, then sum from R6 all the way to R1. You'll find that answer C presents the lowest cost
possible to choose the path R5>R2>R1

upvoted 2 times

Question #244 *Topic 1*

An engineer failed to run diagnostic commands on devices using Cisco DNA Center. Which action in Cisco DNA Center resolves the issue?

A. Enable Secure Shell.

B. Enable APIs.

C. Enable CDP.

D. Enable Command Runner.

Correct Answer: *D*

Community vote distribution

D (100%)

pepgua 6 months, 1 week ago

Selected Answer: D

CHATGPT: Command Runner is a feature available in Cisco DNA Center, a network management platform provided by Cisco. It allows network administrators or engineers to remotely execute commands on multiple network devices simultaneously, providing a centralized and efficient way to manage and configure devices.

Verify Command Runner settings: Confirm that the Command Runner feature is enabled and properly configured in Cisco DNA Center. Ensure that the engineer has the necessary permissions and access rights to use the Command Runner feature.

upvoted 2 times

IceFireSoul 1 year, 2 months ago

Provided answer is correct:
"The Command Runner tool allows you to send diagnostic CLI commands to selected devices."
https://content.cisco.com/chapter.sjs?uri=/searchable/chapter/content/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/1-1/user_guide/b_dnac_ug_1_1/b_dnac_ug_1_1_chapter_01011.html.xml#:~:text=The%20Command%20Runner%20tool%20allows,CLI%20commands%20to%20selected%20devices.

upvoted 3 times

NoUserName1234 1 year, 3 months ago

https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/2-2-2/user_guide/b_cisco_dna_center_ug_2_2_2/b_cisco_dna_center_ug_2_2_2_chapter_0111.html

upvoted 1 times

Question #245 *Topic 1*

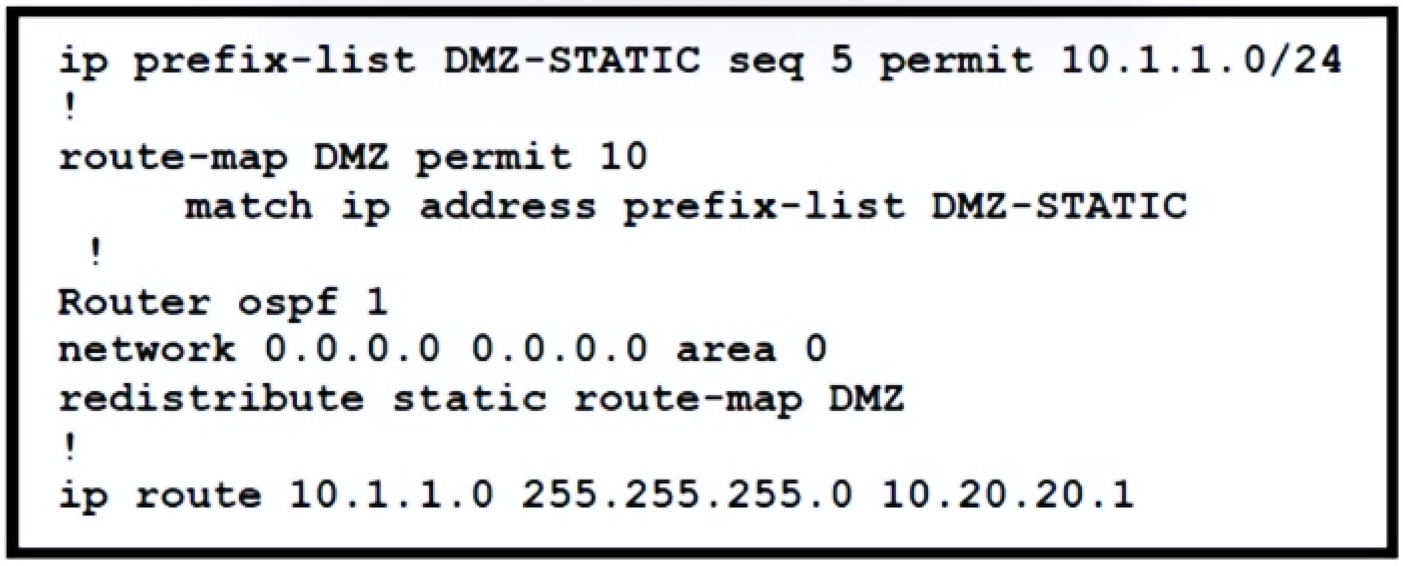

```
ip prefix-list DMZ-STATIC seq 5 permit 10.1.1.0/24
!
route-map DMZ permit 10
    match ip address prefix-list DMZ-STATIC
 !
Router ospf 1
network 0.0.0.0 0.0.0.0 area 0
redistribute static route-map DMZ
!
ip route 10.1.1.0 255.255.255.0 10.20.20.1
```

Refer to the exhibit. The static route is not present in the routing table of an adjacent OSPF neighbor router. Which action resolves the issue?

A. Configure a permit 20 statement to the route map to redistribute the static route.

B. Configure the next-hop interface at the end of the static route for it to get redistributed.

C. Configure the next hop of 10.20.20.1 in the prefix list DMZ-STATIC.

D. Configure the subnets keyword in the redistribution command.

Correct Answer: *D*

Community vote distribution

D (75%) A (25%)

SAMAKEMM 2 months, 2 weeks ago

Selected Answer: D

Without the word "subnets", only the classfull networks will be redistributed

upvoted 1 times

mouin 3 months ago

Selected Answer: D

I tested it in lab without the subnet keyword and without route-map permit 20, the static route got redistributed !!!!so i checked the configuration (show run | sec router ospf) and it turns out that the subnet key word was automatically there

upvoted 1 times

alex711 3 months, 3 weeks ago

D is correct.

https://community.cisco.com/t5/switching/redistribute-static-subnet-to-ospf/td-p/1281958

upvoted 2 times

siyamak 4 months ago

The correct answer is D

Router(config-router)#redistribute sta
Router(config-router)#redistribute static ?
metric Metric for redistributed routes
metric-type OSPF/IS-IS exterior metric type for redistributed routes
nssa-only Limit redistributed routes to NSSA areas
route-map Route map reference
subnets Consider subnets for redistribution into OSPF
tag Set tag for routes redistributed into OSPF
<cr>

upvoted 1 times

Cyril_the_Squirl 4 months, 1 week ago

Selected Answer: A

route-map has a very specific match condition...which is the prefix-list...that is the ONLY thing matched and therefore redistributed. If you want to allow anything else you have to write a condition for it...in this case A is correct.

upvoted 1 times

pepgua 6 months, 1 week ago

Selected Answer: D

```
Router> enable
Router# configure terminal
Router(config)# router ospf <process-id>
Router(config-router)# redistribute static SUBNETS
```

Use the "redistribute" command to redistribute the static route into OSPF. Specify the source of the static route and any necessary parameters.

upvoted 1 times

zhlzjz 9 months, 3 weeks ago

i don't know why? In lab. it is same result without subnets keyword.
all subnets are redistributed whatever classful or not.
Is that different in OS version ?

upvoted 1 times

GReddy2323 9 months, 2 weeks ago

I could be wrong here but I think I have seen videos of redistribution where the instructor states that in latest iOS versions the subnets keywords is already assumed by default and doesn't need to be added. Someone correct me if I am wrong.

upvoted 3 times

HungarianDish 6 months, 3 weeks ago

True. Please see:
https://www.kwtrain.com/blog/route-redistribution-part-1

upvoted 2 times

chris7890 1 year, 2 months ago

the given answer is correct: The command to distribute static route via OSPF in Cisco IOS Router is “redistribute static subnets”
https://www.mustbegeek.com/distribute-static-route-via-ospf-in-cisco-ios-router/

upvoted 2 times

NoUserName1234 1 year, 3 months ago

answer is correct:
https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/2-2-2/user_guide/b_cisco_dna_center_ug_2_2_2/b_cisco_dna_center_ug_2_2_2_chapter_0111.html

upvoted 2 times

Question #246 *Topic 1*

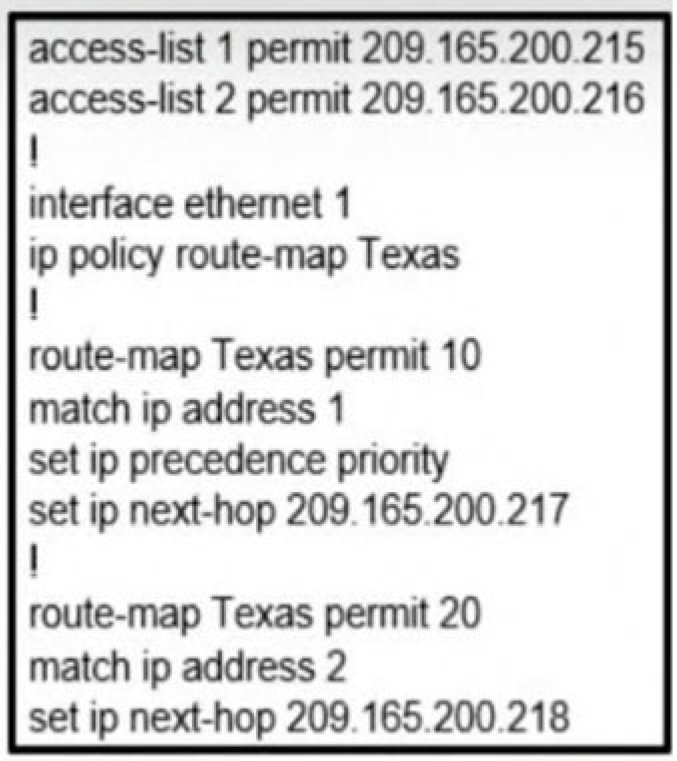

```
access-list 1 permit 209.165.200.215
access-list 2 permit 209.165.200.216
!
interface ethernet 1
ip policy route-map Texas
!
route-map Texas permit 10
match ip address 1
set ip precedence priority
set ip next-hop 209.165.200.217
!
route-map Texas permit 20
match ip address 2
set ip next-hop 209.165.200.218
```

Refer to the exhibit. Packets arriving from source 209.165.200.215 must be sent with the precedence bit set to 1, and packets arriving from source
209.165.200.216 must be sent with the precedence bit set to 5. Which action resolves the issue?

A. set ip precedence critical in route-map Texas permit 20

B. set ip precedence critical in route-map Texas permit 10

C. set ip precedence priority in route-map Texas permit 20

D. set ip precedence immediate in route-map Texas permit 10

Correct Answer: *A*

Community vote distribution

A (100%)

pepgua 6 months, 1 week ago

Selected Answer: A

route-map Texas permit 10
match ip address 1
set ip precedence priority --> bit set to 1 implies priority
set ip next hop x.x.x.x
!!
route-map Texas permit 20
match ip address 2
---------- --> bit set to 5 implies critical (set ip precedence critical)
set ip next hop x.x.x.x

upvoted 1 times

shoo83 11 months, 3 weeks ago

Answer is correct
IP Precedence
000 (0) Routine or Best Effort
001 (1) Priority
010 (2) Immediate
011 (3) Flash - mainly used for Voice Signaling or for Video.
100 (4) Flash Override
101 (5) Critical -mainly used for Voice RTP.
110 (6) Internet
111 (7) Network

upvoted 3 times

jarz 1 year, 2 months ago

Selected Answer: A

Ans is correct

https://www.ccexpert.us/ccie/setting-ip-precedence.html

upvoted 2 times

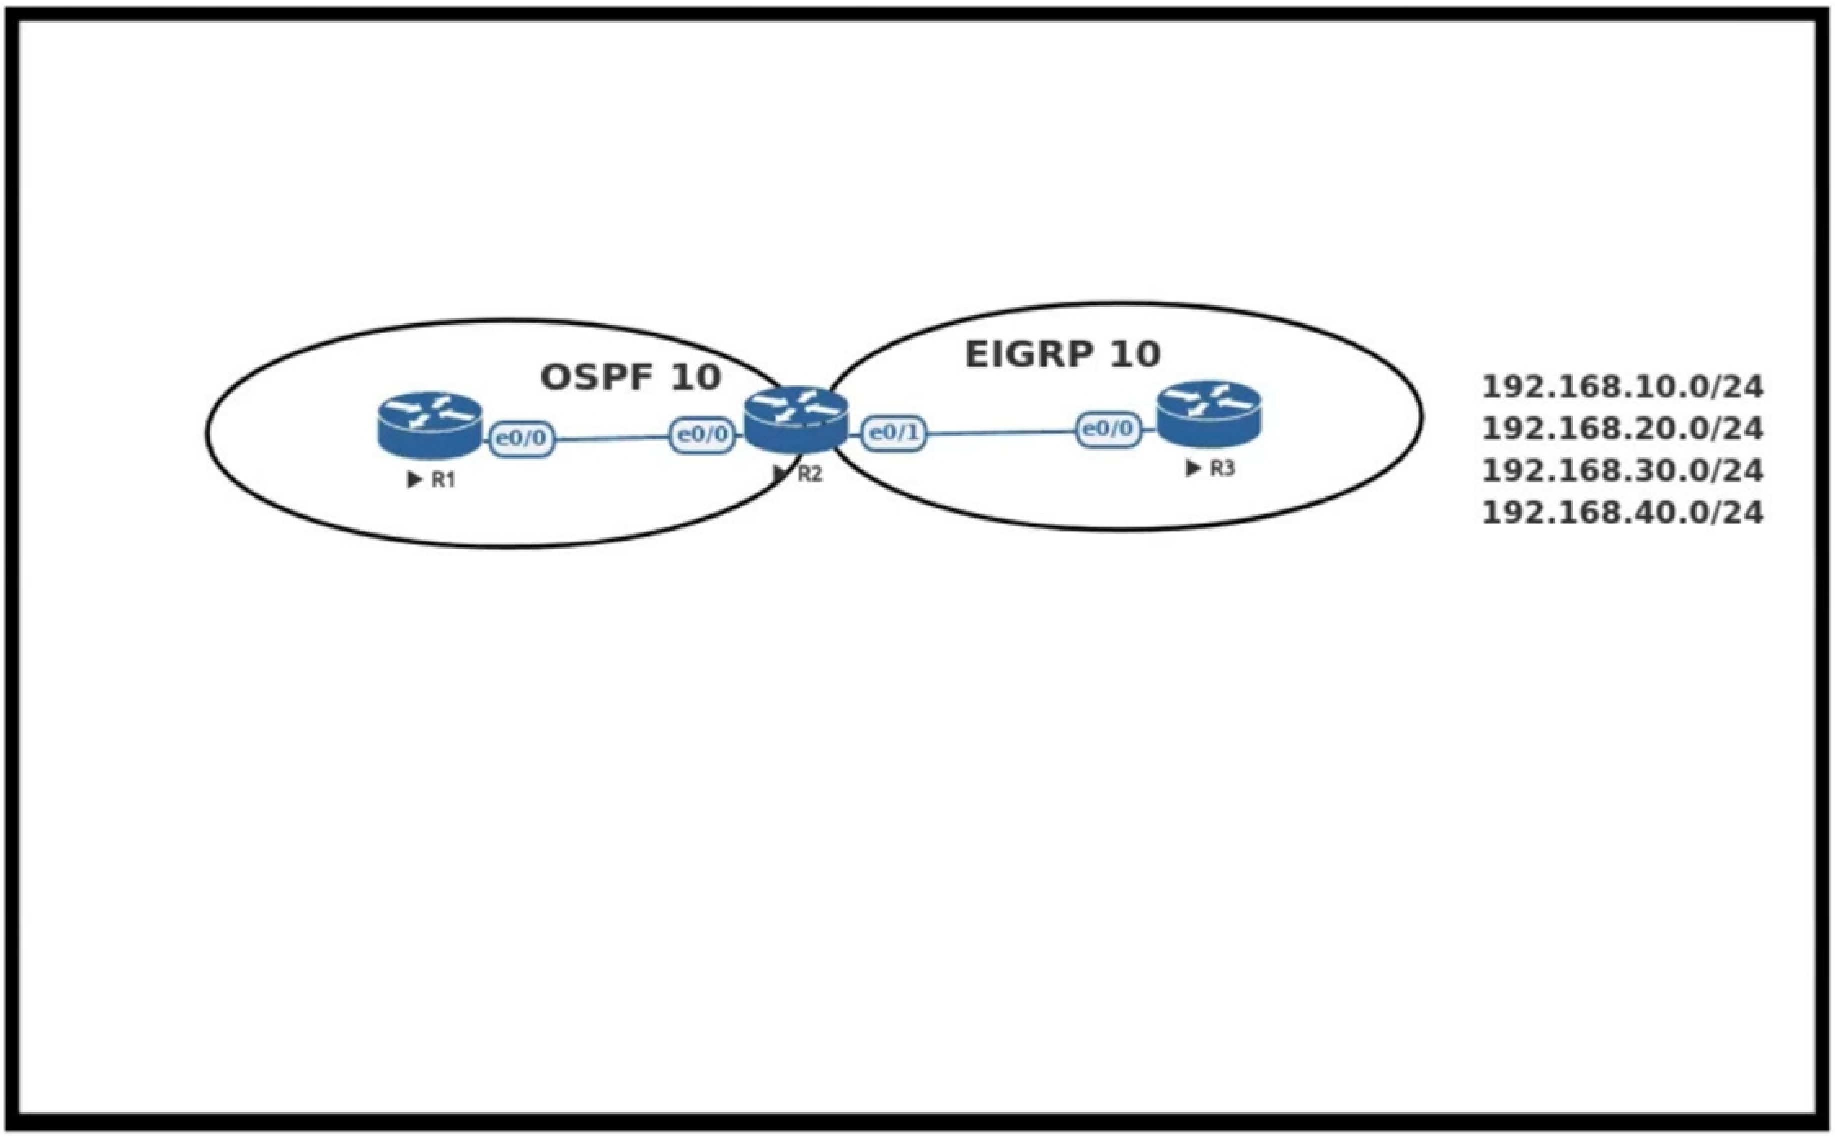

Refer to the exhibit. An engineer must redistribute networks 192.168.10.0/24 and 192.168.20.0/24 into OSPF from EIGRP, where the metric must be added when traversing through multiple hops to start an external route of 20. The engineer notices that the external metric is fixed and does not add at each hop. Which configuration resolves the issue?

A. R2(config)#access-list 10 permit 192.168.10.0 0.0.0.255 R2(config)#access-list 10 permit 192.168.20.0 0.0.0.255 ! R2(config)#route-map RD permit 10 R2(config-route-map)#match ip address 10 R2(config-route-map)#set metric 20 R2(config-route-map)#set metric-type type-2 ! R2(config)#router ospf 10 R2(confjg-router)#redistribute eigrp 10 subnets route-map RD

B. R2(config)#access-list 10 permit 192.168.10.0 0.0.0.255 R2(config)#access-list 10 permit 192.168.20.0 0.0.0.255 ! R2(config)#route-map RD permit 10 R2(config-route-map)#match ip address 10 R2(config-route-map)#set metric 20 R2(config-route-map)#set metric-type type-1 ! R2(config)#router ospf 10 R2(config-router)#redistribute eigrp 10 subnets route-map RD

C. R1(config)#access-list 10 permit 192.168.10.0 0.0.0.255 R1(config)#access-list 10 permit 192.168.20.0 0.0.0.255 ! R1(config)#route-map RD permit 10 R1(config-route-map)#match ip address 10 R1(config-route-map)#set metric 20 R1(config-route-map)#set metric-type type-1 ! R1(config)#router ospf 10 R1(config-router)#redistribute eigrp 10 subnets route-map RD

D. R1(config)#access-list 10 permit 192.168.10.0 0.0.0.255 R1(config)#access-list 10 permit 192.168.20.0 0.0.0.255 ! R1(config)#route-map RD permit 10 R1(config-route-map)#match ip address 10 R1(config-route-map)#set metric 20 R1(config-route-map)#set metric-type type-2 ! R1(config)#router ospf 10 R1(config-router)#redistribute eigrp 10 subnets route-map RD

Correct Answer: *B*

Community vote distribution

B (100%)

chris110 3 months, 1 week ago

B.

```
R2(config)#access-list 10 permit 192.168.10.0 0.0.0.255
R2(config)#access-list 10 permit 192.168.20.0 0.0.0.255
!
R2(config)#route-map RD permit 10
R2(config-route-map)#match ip address 10
R2(config-route-map)#set metric 20
R2(config-route-map)#set metric-type type-1
!
R2(config)#router ospf 10
R2(config-router)#redistribute eigrp 10 subnets route-map RD Most Voted
```

upvoted 1 times

Colmenarez 3 months, 3 weeks ago

Selected Answer: B

R2 & Type1

upvoted 1 times

HungarianDish 7 months, 2 weeks ago

Selected Answer: B

The router connecting OSPF and EIGRP domains by redistribution is an ASBR, in this case it is R2. The configuration needs to be applied there. => Answer "B".

upvoted 1 times

PimplePooper 12 months ago

Selected Answer: B

Given answer is correct. Only external type1 routes will provide the hop count, as it traverses through the network.

upvoted 4 times

Question #248 *Topic 1*

An engineer notices that R1 does not hold enough log messages to identify the root cause during troubleshooting. Which command resolves this issue?

A. #logging buffered 4096 critical

B. #logging buffered 16000 critical

C. (config)#logging buffered 16000 informational

D. (config)#logging buffered 4096 informational

Correct Answer: *C*

Community vote distribution

C (100%)

pepgua 6 months, 1 week ago

Selected Answer: C

(config)# logging buffered <buffer-size>

The default buffer size for log messages in Cisco routers is typically 4096 bytes (4 kilobytes). This default buffer size allows the router to hold a moderate number of log messages before they are overwritten. By increasing the buffer size, it can hold a larger number of log messages, allowing for more extensive logging during troubleshooting.

upvoted 2 times

Slinky 7 months, 4 weeks ago

Answer is C. "logging buffered" command must be entered in the configuration terminal and 4096 is the default size, so we're looking to increase that. the only option here that works is C.

upvoted 4 times

Question #249 *Topic 1*

Which feature minimizes DoS attacks on an IPv6 network?

A. IPv6 Binding Security Table

B. IPv6 Router Advertisement Guard

C. IPv6 Prefix Guard

D. IPv6 Destination Guard

Correct Answer: *D*

Community vote distribution

D (100%)

inteldarvid 5 months, 1 week ago

Selected Answer: D

D correct:

https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/15S/configuration/guide/7600_15_0s_book/IPv6_Security.html#:~:text=on%20the%20VLAN.-,IPv6%20%2D%20Destination%20Guard,%2Dservice%20(DoS)%20attacks.

upvoted 1 times

pepgua 6 months, 1 week ago

Selected Answer: D

From Cisco:
IPv6 - Destination Guard
The Destination Guard feature helps in minimizing denial-of-service (DoS) attacks. It performs address resolutions only for those addresses that are active on the link, and requires the FHS binding table to be populated with the help of the IPv6 snooping feature.

The feature enables the filtering of IPv6 traffic based on the destination address, and blocks the NDP resolution for destination addresses that are not found in the binding table. By default, the policy drops traffic coming for an unknown destination.

upvoted 1 times

NoUserName1234 1 year, 3 months ago

answer is correct ->
https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/15S/configuration/guide/7600_15_0s_book/IPv6_Security.html#86114

upvoted 1 times

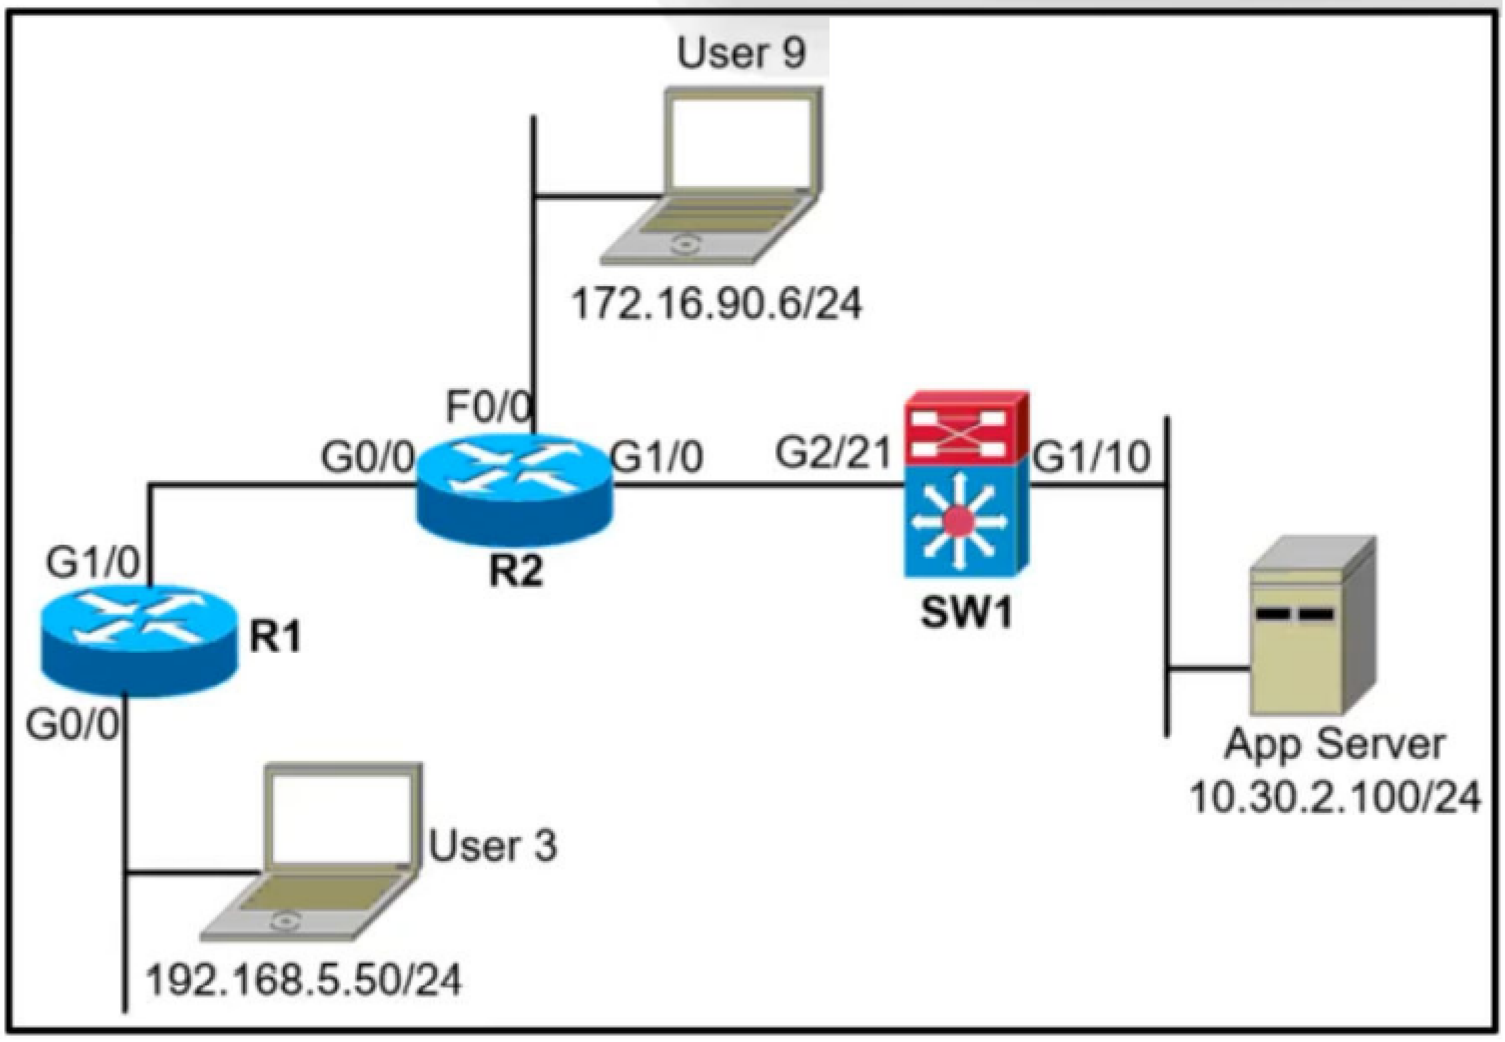

Refer to the exhibit. A network administrator must block ping from user 3 to the App Server only. An inbound standard access list is applied to R1 interface G0/0 to block ping. The network administrator was notified that user 3 cannot even ping user 9 anymore. Where must the access list be applied in the outgoing direction to resolve the issue?

A. R2 interface G0/0

B. SW1 interface G1/10

C. R2 interface G1/0

D. SW1 interface G2/21

Correct Answer: *B*

Community vote distribution

B (58%) C (33%) 8%

Patrick1234 Highly Voted 11 months ago

It's a standard ACL. Standard ACL's should always be installed as close to the DESTINATION as possible. Read this:

Standard ACLs should be located as close to the destination as possible. If a standard ACL were placed at the source of the traffic, the “permit” or “deny” would occur based on the given source address, regardless of the traffic destination.

So the only right answer in this question is B: SW1 interface G1/10.

upvoted 7 times

louisvuitton12 Most Recent 1 month, 2 weeks ago

Selected Answer: B

Closest to the destination

upvoted 1 times

jansan55 2 months, 2 weeks ago

Selected Answer: C

A standard ACL can only deny the IP address of User 3, not only just ping. So the first step to remove that statndard ACL from R1 Gi0/0. We are not sure that SW1 is a an L3 type, so i rule out any SW1 related answers.

upvoted 1 times

Muste 4 months, 1 week ago

Selected Answer: B

provided answer is correct standard access-list should be placed as close to the destination as possible

upvoted 2 times

inteldarvid 5 months, 1 week ago

Selected Answer: D

Correct 100% "D":
team sorry for my earlier reply. The correct answer is "D", it is true, it is the closest to the destination, but it cannot be added (outside or inside) in the swi (g1/10), because the traffic that I want to deny comes from the source and enters the switch through the G2/21, (I tried all the options in my lab) and the correct answer is "D":
SW1 interface G2/21

upvoted 1 times

> **Brand** 4 months, 1 week ago
>
> "Where must the access list be applied in the outgoing direction" It says "outgoing direction" how would you block a traffic sourced by the user3 by applying the ACL to the return traffic back from server?
>
> upvoted 2 times

inteldarvid 5 months, 1 week ago

Selected Answer: B

correct

upvoted 1 times

pepgua 6 months, 1 week ago

Selected Answer: B

By applying the access list in the outgoing direction on the interface facing the App Server, you can ensure that ping traffic from user 3 to other destinations, including user 9, is not affected. Only the ping traffic specifically destined for the App Server will be blocked in the outgoing direction on SW1.

upvoted 1 times

Typovy 8 months, 1 week ago

Selected Answer: B

If vlan's are terminated on switch and then routed to router answer is B. If vlans are terminated on router via .q subinterfaces then answer is C.
Switch icon indicates that this is L3 switch so most propably vlans are ended there on SVI so answet is propably B :)

upvoted 1 times

Jerome_2046 8 months, 2 weeks ago

Selected Answer: B

Standard ACL's should always be installed as close to the DESTINATION as possible

upvoted 1 times

anaisa_goncalves 1 year, 1 month ago

Hi, Why not answer D. Since, it's a standard ACL that has to be applied in outgoing interface. Because if we apply in R2 G1/0, we will not let that User 3 ping SW1, and the question says that it cannot ping ONLY App Server. And this is assuming that SW1 is a layer 3 switch.

upvoted 1 times

> **anaisa_goncalves** 1 year, 1 month ago
>
> Correction I meant option B SW1 Interface Gi 1/10 as correct answer
>
> upvoted 2 times

VergilP 1 year, 1 month ago

I am confuse of question is ask about..
so question is ask ..delete R1 G0/0 ACL and place the ACL "somewhere" then make User3 can ping User9 but can not reach app server?
Is my understanding correct?

upvoted 1 times

Remsync 1 year, 2 months ago

Selected Answer: C

If you're usign an ACL to block ping, that means you're using an extended ACL, and it's recommended to place de ACL closest to the source, so the given answer is correct.

upvoted 1 times

> **Remsync** 1 year, 2 months ago
>
> My bad, it says standard ACL. Given answer is correct.
>
> upvoted 2 times

> **Remsync** 1 year, 2 months ago
>
> I mean, C is correct, not the given answer.
>
> upvoted 1 times

IceFireSoul 1 year, 2 months ago

Given answer is correct, at least by my standards. Switch device in this diagram is not a pure layer 2 switch, in fact its a layer 3 switch and therefor can make routing decisions as well , in this case block ping going out the interface G1/10

upvoted 3 times

> **Remsync** 1 year, 2 months ago
>
> My bad, it says standard ACL. Given answer is correct. You're correct.
>
> upvoted 1 times

Remsync 1 year, 2 months ago

If you're usign an ACL to block ping, that means you're using an extended ACL, and it's recommended to place de ACL closest to the source, so the given answer is correct.

By putting the ACL on the L3 SW it goes against that principle since you're placing it closes to the destination.

upvoted 1 times

Remsync 1 year, 2 months ago

I mean, C is correct, not the given answer.

upvoted 1 times

chris7890 1 year, 3 months ago

Selected Answer: C

Where must the access list be applied in the outgoing direction to resolve the issue?
Answer C must be correct!

upvoted 2 times

lisanta12 1 year, 3 months ago

No, in the case of C, ping cannot be executed until SW1.

upvoted 1 times

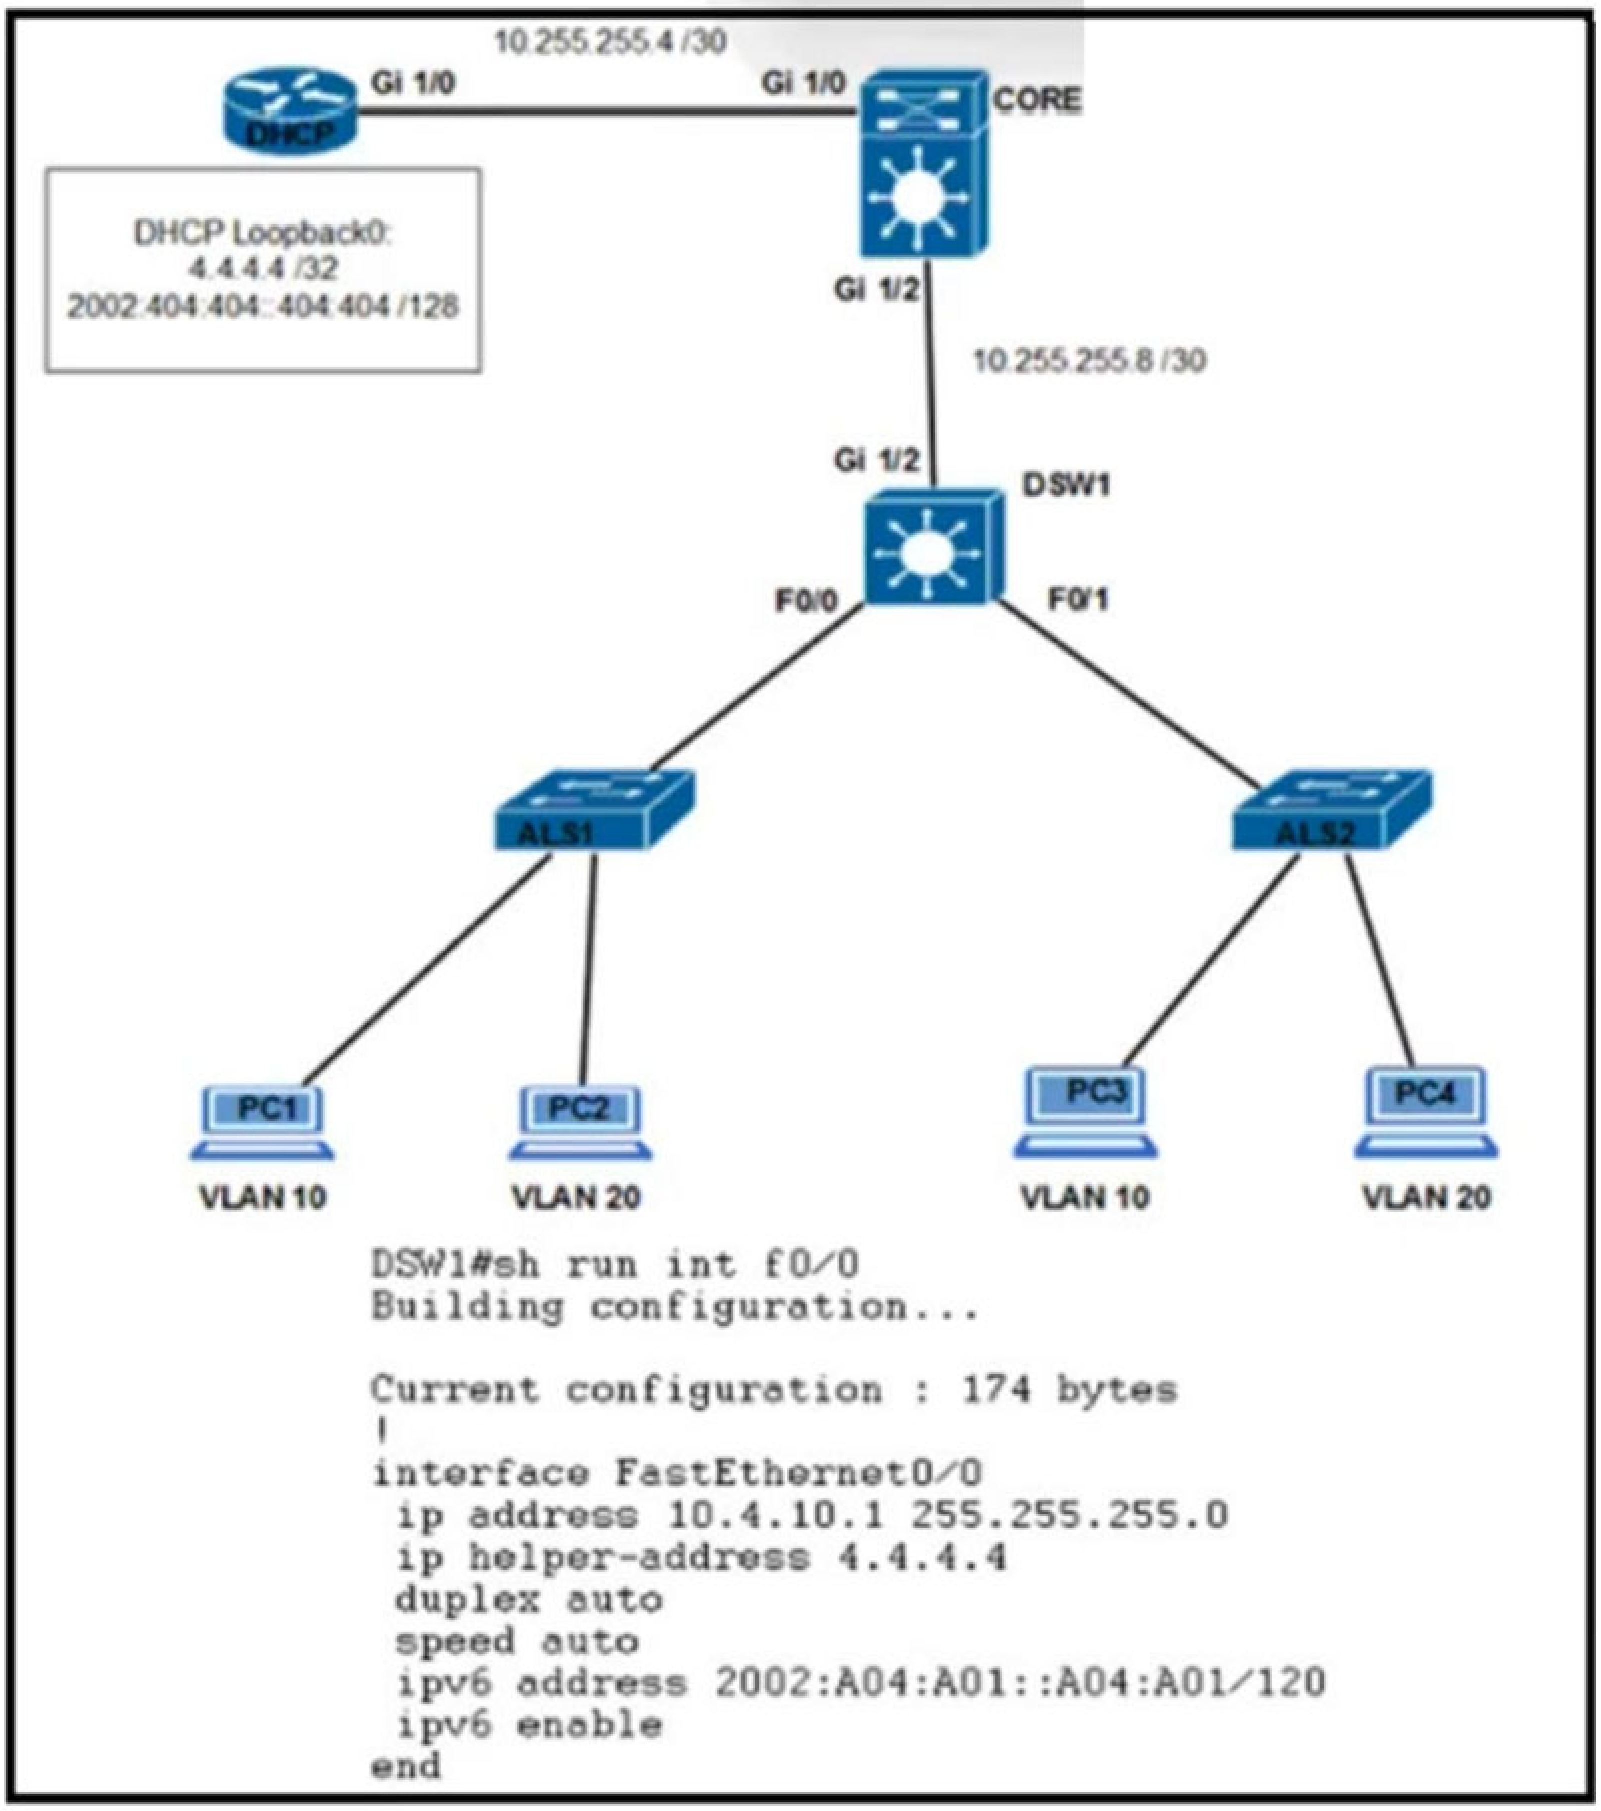

```
DSW1#sh run int f0/0
Building configuration...

Current configuration : 174 bytes
!
interface FastEthernet0/0
 ip address 10.4.10.1 255.255.255.0
 ip helper-address 4.4.4.4
 duplex auto
 speed auto
 ipv6 address 2002:A04:A01::A04:A01/120
 ipv6 enable
end
```

Refer to the exhibit. Clients on ALS2 receive IPv4 and IPv6 addresses, but clients on ALS1 receive only IPv4 addresses and not IPv6 addresses. Which action on
DSW1 allows clients on ALS1 to receive IPv6 addresses?

A. Configure DSW1(dhcp-config)#default-router 2002:A04:A01::A04:A01

B. Configure DSW1(config-if)#ipv6 dhcp relay destination 2002:404:404::404:404 GigabitEthernet1/2

C. Configure DSW1(config)#ipv6 route 2002:404:404::404:404/128 FastEthernet1/0

D. Configure DSW1(config-if)#ipv6 helper address 2002:404:404::404:404

Correct Answer: *B*

Community vote distribution

B (100%)

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: B

B is correct

upvoted 1 times

⊟ **slcc99** 6 months, 1 week ago

Why do ALS2 clients receive IPv6 addresses?
upvoted 1 times

Rob_CCNP000 6 months ago

DSW1 Fa0/0 and Fa0/1 are layer 3 interfaces so both need the dhcp relay configured.
upvoted 1 times

HungarianDish 7 months, 2 weeks ago

Selected Answer: B

```
DSW1(config)#int f0/0
DSW1(config-if)#ipv6 dhcp relay destination 2002:404:404::404:404 GigabitEthernet1/2
```

Explanation:
https://www.cbtnuggets.com/blog/technology/networking/how-to-use-the-ipv6-dhcp-relay
upvoted 4 times

Question #252 Topic 1

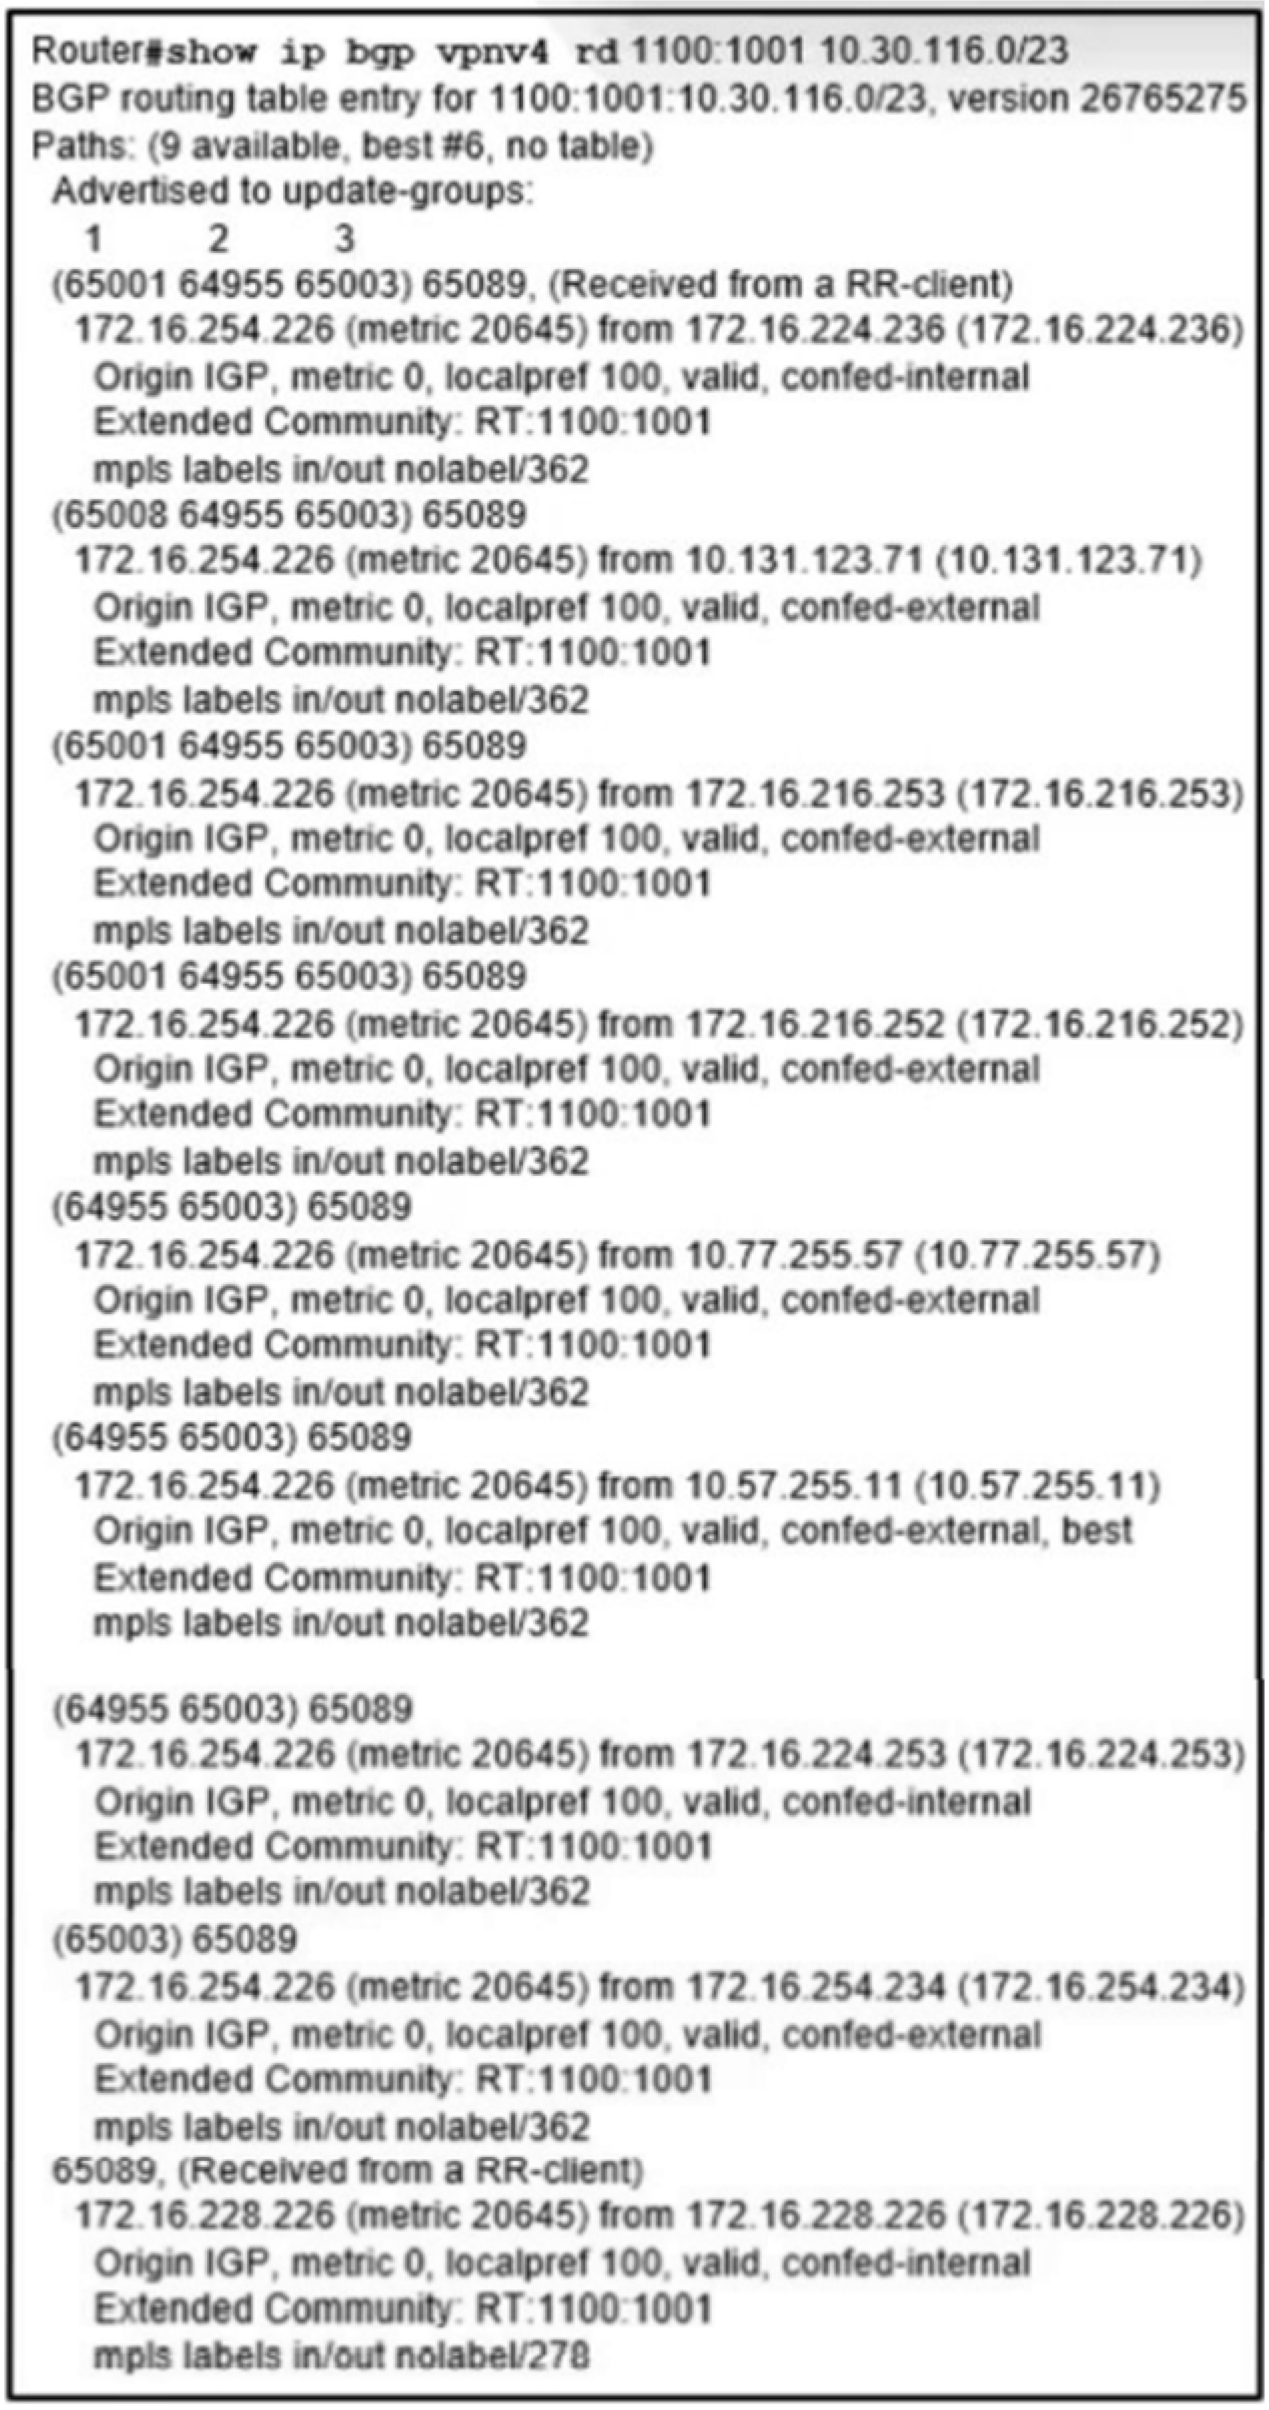

```
Router#show ip bgp vpnv4 rd 1100:1001 10.30.116.0/23
BGP routing table entry for 1100:1001:10.30.116.0/23, version 26765275
Paths: (9 available, best #6, no table)
 Advertised to update-groups:
   1        2        3
 (65001 64955 65003) 65089, (Received from a RR-client)
  172.16.254.226 (metric 20645) from 172.16.224.236 (172.16.224.236)
    Origin IGP, metric 0, localpref 100, valid, confed-internal
    Extended Community: RT:1100:1001
    mpls labels in/out nolabel/362
 (65008 64955 65003) 65089
  172.16.254.226 (metric 20645) from 10.131.123.71 (10.131.123.71)
    Origin IGP, metric 0, localpref 100, valid, confed-external
    Extended Community: RT:1100:1001
    mpls labels in/out nolabel/362
 (65001 64955 65003) 65089
  172.16.254.226 (metric 20645) from 172.16.216.253 (172.16.216.253)
    Origin IGP, metric 0, localpref 100, valid, confed-external
    Extended Community: RT:1100:1001
    mpls labels in/out nolabel/362
 (65001 64955 65003) 65089
  172.16.254.226 (metric 20645) from 172.16.216.252 (172.16.216.252)
    Origin IGP, metric 0, localpref 100, valid, confed-external
    Extended Community: RT:1100:1001
    mpls labels in/out nolabel/362
 (64955 65003) 65089
  172.16.254.226 (metric 20645) from 10.77.255.57 (10.77.255.57)
    Origin IGP, metric 0, localpref 100, valid, confed-external
    Extended Community: RT:1100:1001
    mpls labels in/out nolabel/362
 (64955 65003) 65089
  172.16.254.226 (metric 20645) from 10.57.255.11 (10.57.255.11)
    Origin IGP, metric 0, localpref 100, valid, confed-external, best
    Extended Community: RT:1100:1001
    mpls labels in/out nolabel/362

 (64955 65003) 65089
  172.16.254.226 (metric 20645) from 172.16.224.253 (172.16.224.253)
    Origin IGP, metric 0, localpref 100, valid, confed-internal
    Extended Community: RT:1100:1001
    mpls labels in/out nolabel/362
 (65003) 65089
  172.16.254.226 (metric 20645) from 172.16.254.234 (172.16.254.234)
    Origin IGP, metric 0, localpref 100, valid, confed-external
    Extended Community: RT:1100:1001
    mpls labels in/out nolabel/362
 65089, (Received from a RR-client)
  172.16.228.226 (metric 20645) from 172.16.228.226 (172.16.228.226)
    Origin IGP, metric 0, localpref 100, valid, confed-internal
    Extended Community: RT:1100:1001
    mpls labels in/out nolabel/278
```

Refer to the exhibit. An engineer configured BGP and wants to select the path from 10.77.255.57 as the best path instead of current best path. Which action resolves the issue?

A. Configure higher MED to select as the best path.

B. Configure AS_PATH prepend for the current best path.

C. Configure AS_PATH prepend for the desired best path.

D. Configure lower LOCAL_PREF to select as the best path.

Correct Answer: *B*

Community vote distribution

B (100%)

Zizu007 Highly Voted 11 months, 3 weeks ago

Selected Answer: B

Output shows #9 different possible paths. local routers has chosen #6 as best-path. it is asked to what is needed to make path #5 the best-path.
- A - wrong (lower MED is preferred.)
- B - correct (by adding extra AS_PATH makes the current best-path #6 less preferred compared to route #5)
- C - wrong (this is the opposite of B)
- D - wrong (higher LOCAL_PREF is preferred not lower!)

upvoted 6 times

ZamanR Most Recent 4 days, 23 hours ago

D is correct i think
Explanation

From the output, we learn that the current best path is from 10.57.255.11 (which includes "...valid,

confed-external, best") and this path is 2 ASes away (64955 65003). Although there are some paths

with only 1 AS away (path from 172.16.254.234 for example) but they were not chosen the best path

so AS_PATH was not used to determine the best path -> Answers A and answer C are not correct.

All the paths in the output have metric of 0 and this is the lowest (best) value for this attribute. If we

configure higher MED then it is less preferred over other paths -> Answer B is not correct.

Only answer D is left but LOCAL_PREF attribute should be configured with higher value to be preferred

so we hope "lower LOCAL_PREF" here means higher value. But this is the best answer

upvoted 1 times

HungarianDish 7 months, 2 weeks ago

Selected Answer: B

My assumption is that the best path is chosen for the lowest BGP router-id, the lowest is 10.57.255.11 and the second lowest is 10.77.255.57.
If we make 10.57.255.11 less preferred by AS Path Prepending, 10.77.255.57 is going to be selected as best.
All other attributes are the same.

upvoted 1 times

ClaudeYun 8 months, 2 weeks ago

Although B is sort of making sense and according to commen sense, it still hard to convince answer B is correct due to there's other BGP routes with less AS path and the same other attributes.
e.g. metric, localpref but not be choosen the best.
E.g. 172.16.254.234. it's a tricky question.

upvoted 2 times

Jerome_2046 8 months, 2 weeks ago

From the output, we learn that the current best path is from 10.57.255.11 (which includes "...valid, confed-external, best") and this path is 2 ASes away (64955 65003).
Although there are some paths with only 1 AS away (path from 172.16.254.234 for example) but they were not chosen the best path, so AS_PATH was not used to determine the best path. Answers A and answer C are not correct.

upvoted 1 times

babs 12 months ago

can someone explain

upvoted 1 times

Zizu007 11 months, 3 weeks ago

Output shows #9 different possible paths. local routers has chosen #6 as best-path. it is asked to what is needed to make path #5 the best-path.
- A - wrong (lower MED is preferred.)
- B - correct (by adding extra AS_PATH makes the current best-path #6 less preferred compared to route #5)
- C - wrong (this is the opposite of B)
- D - wrong (higher LOCAL_PREF is preferred not lower!)

upvoted 2 times

Question #253 *Topic 1*

What is a function of IPv6 Source Guard?

A. It inspects ND and DHCP packets to build an address binding table.

B. It works with address glean or ND to find existing addresses.

C. It notifies the ND protocol to inform hosts if the traffic is denied by it.

D. It denies traffic from known sources and allocated addresses.

Correct Answer: *B*

Community vote distribution

B (67%) A (22%) 11%

NoUserName1234 Highly Voted 1 year, 2 months ago

Selected Answer: B

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-3s/ip6f-xe-3s-book/ip6-src-guard.html

upvoted 5 times

JKStinn 1 year ago

IPv6 source guard does not inspect ND or DHCP packets; rather, it works in conjunction with IPv6 neighbor discovery (ND) inspection or IPv6 address glean, both of which detect existing addresses on the link and store them into the binding table.

upvoted 2 times

Hermin 9 months, 3 weeks ago

Source Guard only looks at information found in the binding table, and it doesn't fill the binding table. You need another feature like ND inspection or IPv6 snooping to do this.
https://networklessons.com/cisco/ccie-routing-switching-written/ipv6-source-guard

upvoted 2 times

inteldarvid Most Recent 5 months, 1 week ago

Selected Answer: B

very sorry team, with my question before, the option correct is ""B"", look this info:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-16/ip6f-xe-16-book/ip6-src-guard.pdf

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: D

team is option D:
IPv6 source guard can deny traffic from unknown sources or unallocated addresses, such as traffic from sources not assigned by a DHCP server. When traffic is denied, the IPv6 address glean feature is notified so that it can try to recover the traffic by querying the DHCP server or by using IPv6 ND

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/15-e/ip6f-15-e-book/ip6f-15-e-book_chapter_0110.pdf

upvoted 1 times

jarz 1 year, 2 months ago

Selected Answer: A

An entry is installed in the binding table when one of the following conditions is satisfied:
• An IPv6 binding is learnt through DHCP.
• An IPv6 address or prefix is learnt through NDP.
• A static binding is configured by the user.

Source
https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/15S/configuration/guide/7600_15_0s_book/IPv6_Security.pdf

upvoted 2 times

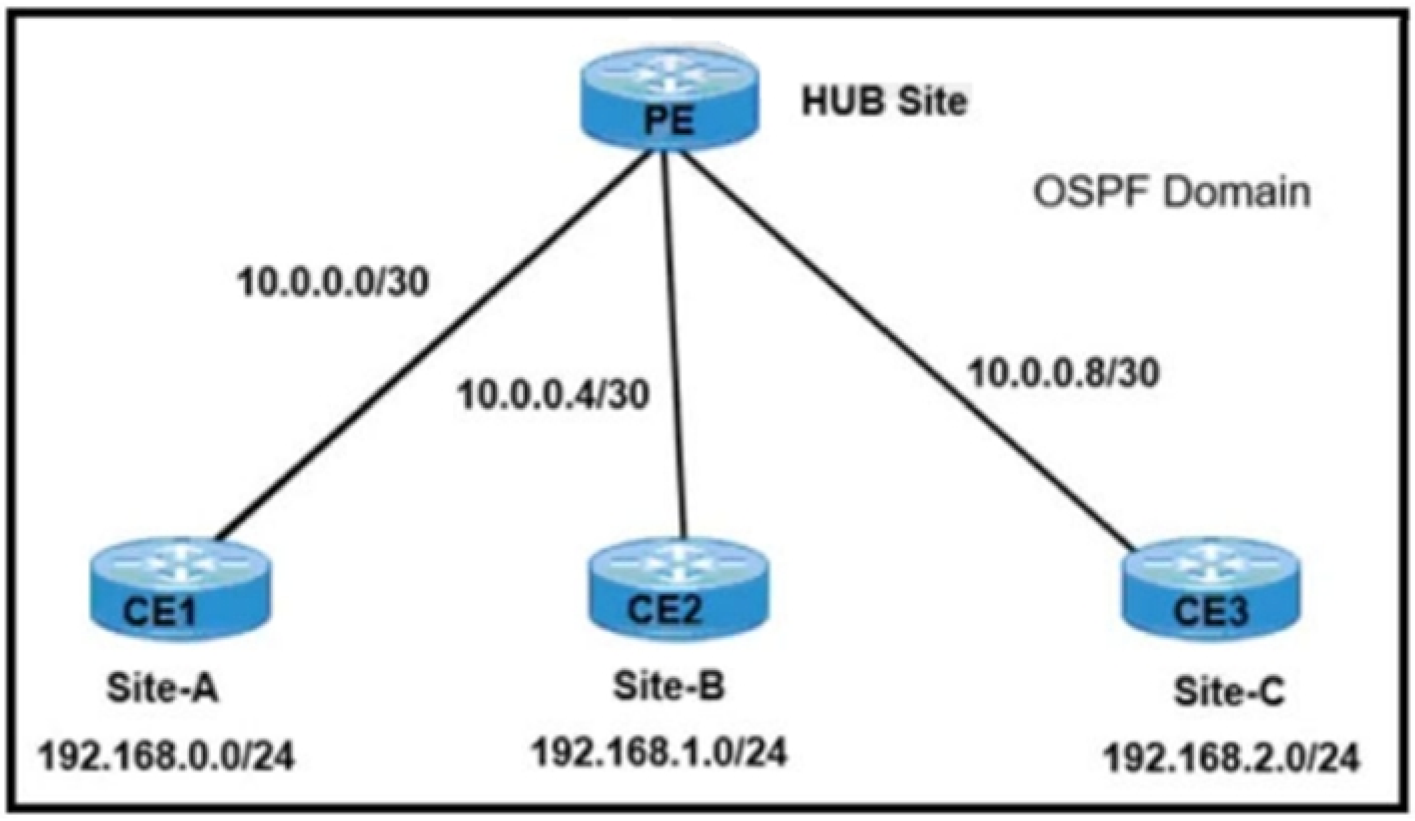

Refer to the exhibit. A network engineer must establish communication between three different customer sites with these requirements:

* Site-A: must be restricted to access to any users at Site-B or Site-C.

* Site-B and Site-C: must be able to communicate between sites and share routes using OSPF.

PE interface configuration:

```
interface FastEthernet0/0
ip vrf forwarding Site-A
!
interface FastEthernet0/1
ip vrf forwarding SharedSites
!
interface FastEthernet0/2
ip vrf forwarding SharedSites
```

Which configuration meets the requirements?

A. PE(config)#router ospf 10 vrf Site-A PE(config-router)#network 0.0.0.0 255.255.255.255 area 0 PE(config)#router ospf 10 vrf SharedSites PE(config-router)#network 0.0.0.0 255.255.255.255 area 1

B. PE(config)#router ospf 10 vrf Site-A PE(config-router)#network 0.0.0.0 255.255.255.255 area 0 PE(config)#router ospf 20 vrf SharedSites PE(config-router)#network 0.0.0.0 255.255.255.255 area 1

C. PE(config)#router ospf 10 vrf Site-A PE(config-router)#network 0.0.0.0 255.255.255.255 area 0 PE(config)#router ospf 10 vrf SharedSites PE(config-router)#network 0.0.0.0 255.255.255.255 area 0

D. PE(config)#router ospf 10 vrf Site-A PE(config-router)#network 0.0.0.0 255.255.255.255 area 0 PE(config)#router ospf 20 vrf SharedSites PE(config-router)#network 0.0.0.0 255.255.255.255 area 0

Correct Answer: *C*

Community vote distribution

D (100%)

⊟ **jarz** Highly Voted 1 year, 2 months ago

Selected Answer: D

Answer is D
And before you ask, you need unique process IDs in each VRF.

upvoted 12 times

⊟ **inteldarvid** Most Recent 5 months, 1 week ago

Selected Answer: D

Correct is "D"

upvoted 1 times

⊟ **chris7890** 1 year, 3 months ago

Selected Answer: D

I think the correct Answer is D

upvoted 4 times

IceFireSoul 1 year, 2 months ago

Why you think answer is D ? What your reasoning ?

upvoted 1 times

Almylle 5 months, 4 weeks ago

Because you need to configure diferent process ID, the area is independent because of the VRF, but in this case the process ID is not independent.

upvoted 1 times

jarz 1 year, 2 months ago

you need to configure different process ID in the VRFs

upvoted 3 times

Question #255 *Topic 1*

What is LDP label binding?

A. destination prefix with label

B. two routers with label distribution session

C. source prefix with label

D. neighboring router with label

Correct Answer: *A*

Reference:

https://www.cisco.com/en/US/docs/general/Test/kwoodwar/fsinbd4.pdf

Community vote distribution

A (100%)

Malasxd 7 months, 2 weeks ago

"A"

label binding—An association between a destination prefix and a label.

https://www.cisco.com/en/US/docs/general/Test/kwoodwar/fsinbd4.pdf

upvoted 1 times

HungarianDish 7 months, 2 weeks ago

Selected Answer: A

"label binding—An association between a destination prefix and a label."

https://www.cisco.com/en/US/docs/general/Test/kwoodwar/fsinbd4.pdf

Source and answer are correct

upvoted 1 times

Question #256 *Topic 1*

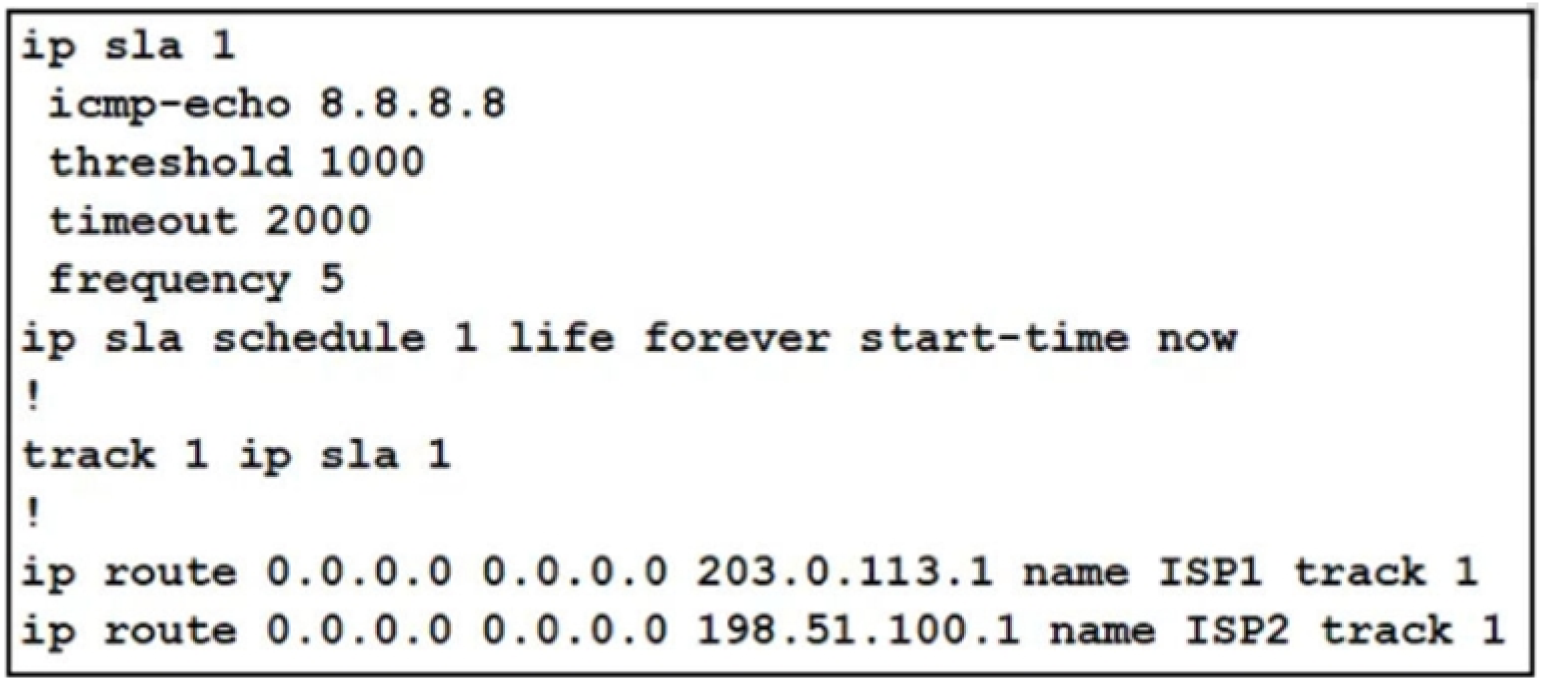

```
ip sla 1
 icmp-echo 8.8.8.8
 threshold 1000
 timeout 2000
 frequency 5
ip sla schedule 1 life forever start-time now
!
track 1 ip sla 1
!
ip route 0.0.0.0 0.0.0.0 203.0.113.1 name ISP1 track 1
ip route 0.0.0.0 0.0.0.0 198.51.100.1 name ISP2 track 1
```

Refer to the exhibit. An administrator configures a router to stop using a particular default route if the DNS server 8.8.8.8 is not reachable through that route.

However, this configuration did not work as desired and the default route still works even if the DNS server 8.8.8.8 is unreachable. Which two configuration changes resolve the issue? (Choose two.)

A. Use a separate track object to reference the existing IP SLA 1 probe for every static route.

B. Use a separate IP SLA probe and track object for every static route.

C. Associate every IP SLA probe with the proper WAN address of the router.

D. Reference the proper exit interfaces along with the next hops in both static default routes.

E. Configure two static routes for the 8.8.8.8/32 destination to match the IP SLA probe for each ISP.

Correct Answer: *BE*

Community vote distribution

BC (100%)

Almylle 5 months, 4 weeks ago

Selected Answer: BC

Im going for B and C, because the alternative E it's isn't needed, with the default route u already have communication with google DNS, so you only need to separate the tracks between static routes and WAN's

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

It is not clearly described what they want to configure with "E". It could be a valid option with the correct configuration. Based on this:
https://community.cisco.com/t5/routing/ip-sla-tracking-a-far-ip/td-p/1971337

The first two static routes are there to make sure that the tested IP address inside the ISP1 is truly reached only via link to ISP1,
and if that link is down, then the pings are going to be thrown away (this is to prevent pinging 8.8.8.8 via ISP2 thanks to the default route).

```
ip route 8.8.8.8 255.255.255.255 Ethernet0/0 10.0.0.1
ip route 8.8.8.8 255.255.255.255 Null0 2
ip sla 1
icmp-echo 8.8.8.8 source-interface Ethernet0/0
threshold 800
timeout 1000
frequency 30
ip sla schedule 1 start-time now life forever
track 1 rtr 1 reachability
ip route 0.0.0.0 0.0.0.0 10.0.0.1 track 1
ip route 0.0.0.0 0.0.0.0 20.0.0.1 2
```

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

This way, a static route to 8.8.8.8/32 should be set only via primary ISP. Not for both ISPs. This makes "E" incorrect.

upvoted 1 times

Lilienen 10 months, 2 weeks ago

Selected Answer: BC

Correct:
B - Both static routes must have a separate Track object and IP SLA probe.
C - Each SLA probe must originate from a different ISP, therefore a different IP address.

Wrong:
A - Only a separate Track object won't help with anything, we need also a separate IP SLA probe.
D - This is redundant, the router knows which interface to use for both next hops (based on ARP and MAC address table).
E - This is just messy and not needed, we just need to set a different source for each probe (answer C).

upvoted 3 times

ellen_AA 11 months, 2 weeks ago

D & E are correct

upvoted 2 times

Huntkey 1 year, 2 months ago

Selected Answer: BC

I would vote for B and C. Setting the static route to 8.8.8.8 for both ISP doesn't make sense. It would make sense if it is only one static route for that.

upvoted 2 times

NoUserName1234 1 year, 2 months ago

https://community.cisco.com/t5/routing/ip-sla-tracking-a-far-ip/td-p/1971337

upvoted 2 times

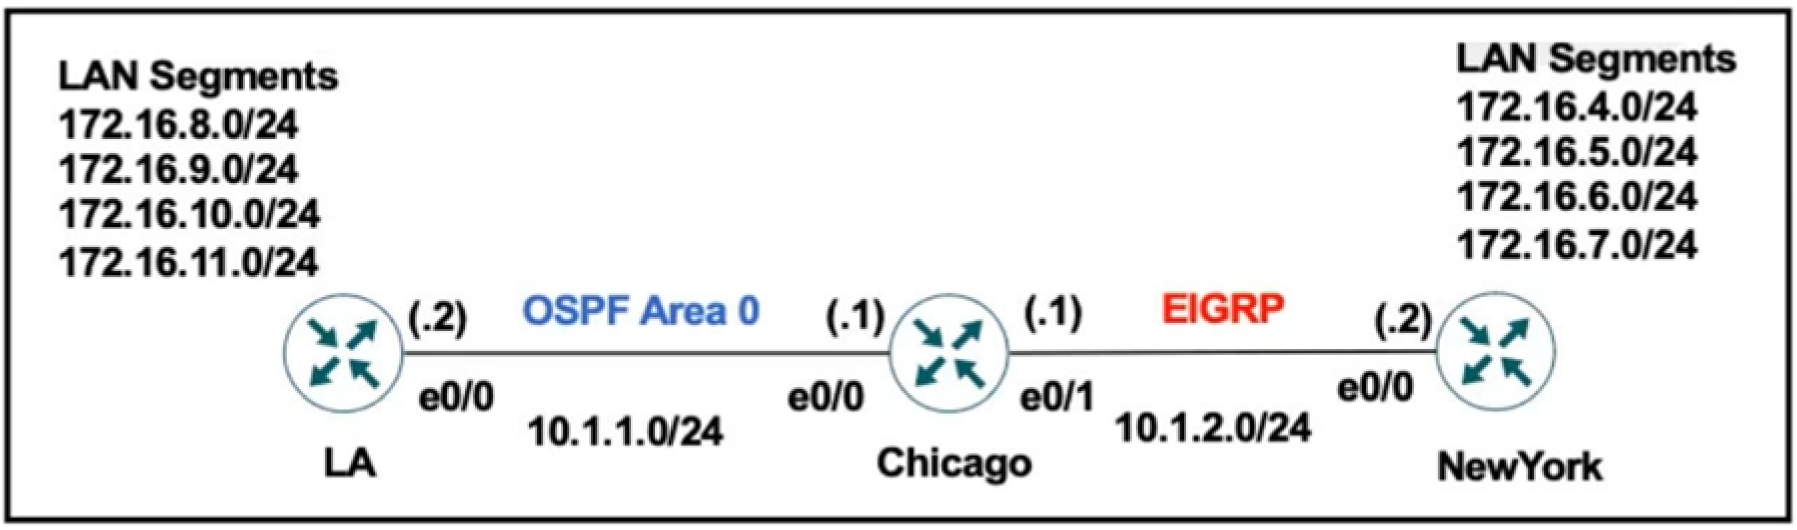

Refer to the exhibit. The network administrator configured the Chicago router to mutually redistribute the LA and NewYork routes with OSPF routes to be summarized as a single route in EIGRP using the longest summary mask: router eigrp 100 redistribute ospf 1 metric 10 10 10 10 10 router ospf 1 redistribute eigrp 100 subnets

!

interface E 0/0

ip summary-address eigrp 100 172.16.0.0 255.255.0.0

After the configuration, the New York router receives all the specific LA routes but the summary route. Which set of configurations resolves the issue on the

Chicago router?

A. router eigrp 100 summary-address 172.16.8.0 255.255.252.0

B. interface E 0/1 ip summary-address eigrp 100 172.16.8.0 255.255.252.0

C. router eigrp 100 summary-address 172.16.0.0 255.255.0.0

D. interface E 0/1 ip summary-address eigrp 100 172.16.0.0 255.255.0.0

Correct Answer: *B*

Community vote distribution

B (88%) 13%

ChillingAgain Highly Voted 1 year, 1 month ago

Selected Answer: B

Answer is B. Summarized route for 172.16.8.0/24, 172.16.9.0/24, 172.16.10.0/24, 172.16.11.0/24 is 172.16.8.0/22. Which is noted as 172.16.8.0 255.255.252.0

upvoted 6 times

inteldarvid Most Recent 5 months, 1 week ago

Selected Answer: B

B correct:

The advertisement of summary routes occurs on an interface-by-interface basis. For classic EIGRP configuration mode, you use the interface parameter command ip summary-address eigrp as-number network subnet-mask [leak-map route-map-name] to place an EIGRP summary aggregate on an interface.

upvoted 1 times

Remsync 1 year, 1 month ago

Selected Answer: D

D is the correct answer. Even though both B and D solve the problem, the question is asking for the longest summary mask. D is summarizing with a /16 while B is doing it with a /22. /16 is longer than a /22. Answer D.

upvoted 1 times

ChillingAgain 1 year, 1 month ago

Longer prefix means more subnet bits. So /22 is longer than /16.
So answer is B

upvoted 10 times

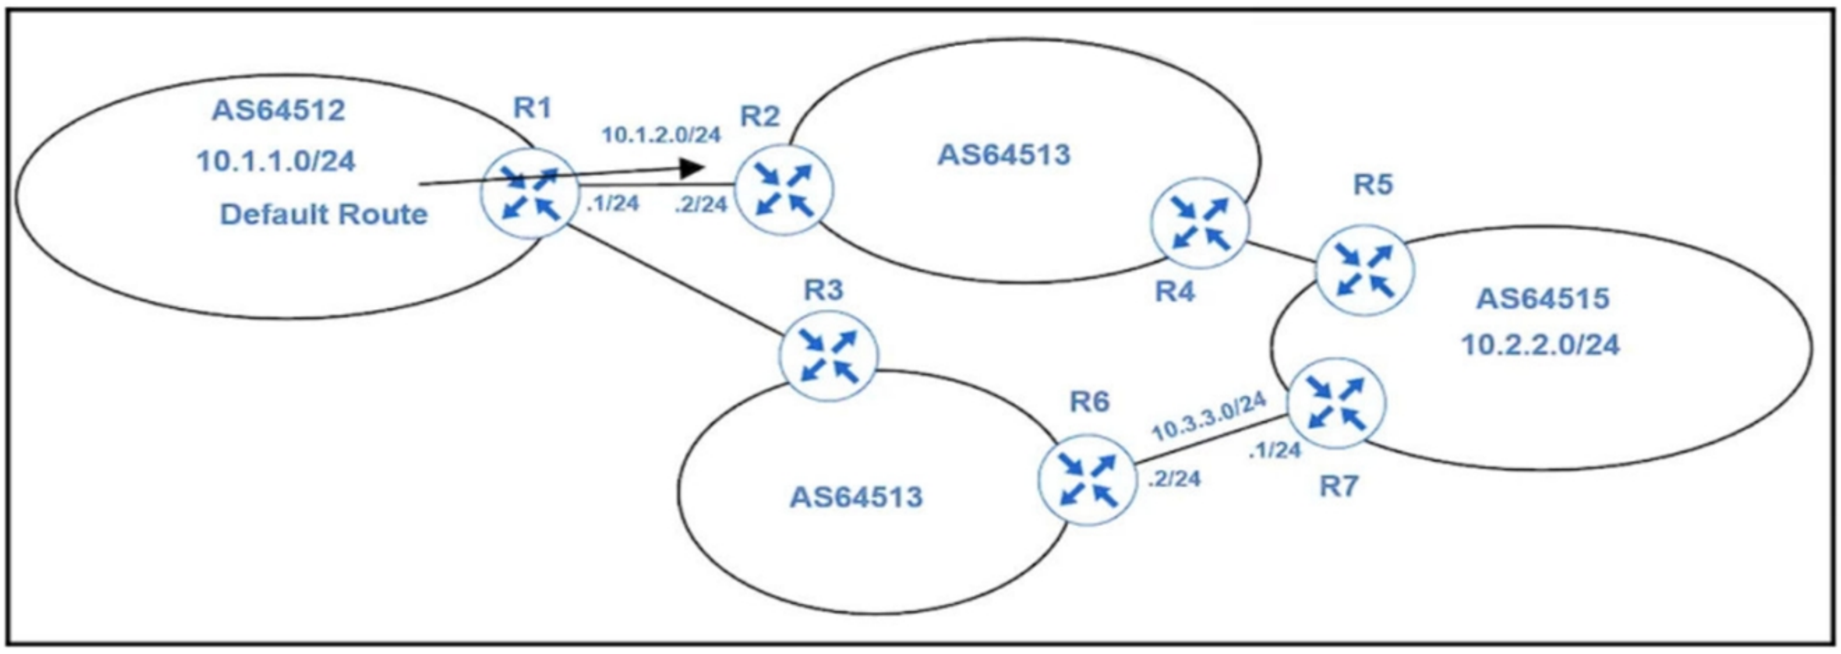

Refer to the exhibit. An engineer must configure PBR on R1 to reach to 10.2.2.0/24 via R3 AS64513 as the primary path and a backup route through default route via R2 AS64513. All BGP routes are in the routing table of R1, but a static default route overrides BGP routes. Which PBR configuration achieves the objective?

A. access-list 100 permit ip 10.1.1.0 0.0.0.255 10.2.2.0 0.0.0.255 ! route-map PBR permit 10 match ip address 100 set ip next-hop recursive 10.3.3.1

B. access-list 100 permit ip 10.1.1.0 255.255.255.0 10.2.2.0 255.255.255.0 ! route-map PBR permit 10 match ip address 100 set ip next-hop recursive 10.3.3.1

C. access-list 100 permit ip 10.1.1.0 0.0.0.255 10.2.2.0 0.0.0.255 ! route-map PBR permit 10 match ip address 100 set ip next-hop 10.3.3.1

D. access-list 100 permit ip 10.1.1.0 255.255.255.0 10.2.2.0 255.255.255.0 ! route-map PBR permit 10 match ip address 100 set ip next-hop 10.3.3.1

Correct Answer: *A*

Community vote distribution

A (100%)

⊟ **shoo83** Highly Voted 11 months, 3 weeks ago

Answer is correct (A)
The PBR Recursive Next Hop feature enhances route maps to enable configuration of a recursive next-hop IP address that is used by policy-based routing (PBR). The recursive next-hop IP address is installed in the routing table and can be a subnet that is not directly connected. If the recursive next-hop IP address is not available, packets are routed using a default route.

upvoted 5 times

⊟ **[Removed]** Most Recent 4 months, 4 weeks ago

Am I blind or is answer A and C the same?

upvoted 2 times

⊟ **[Removed]** 4 months, 4 weeks ago

Disregard, I missed the keyword recursive under A

upvoted 1 times

⊟ **Colmenarez** 3 months, 3 weeks ago

Spot the difference type of question

upvoted 1 times

⊟ **inteldarvid** 5 months, 1 week ago

A correct:
https://notes.networklessons.com/pbr-next-hop-recursive

upvoted 1 times

⊟ **Aikat** 9 months, 2 weeks ago

Selected Answer: A

Answer is C

The PBR Recursive Next Hop feature enhances route maps to enable configuration of a recursive next-hop IP address that is used by policy-based routing (PBR). The recursive next-hop IP address is installed in the routing table and can be a subnet that is not directly connected. In this case 10.3.3.1 is a subnet which is not directly connected.

upvoted 2 times

Aikat 9 months, 2 weeks ago

I meant: Answer is A

upvoted 2 times

chris7890 11 months, 3 weeks ago

Is this answer correct? As the Cisco document states: Note
This configuration does not ensure that packets get routed using the recursive IP address if an intermediate IP address is a shorter route to the destination.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/xe-3s/iri-xe-3s-book/iri-pbr-rec-next-hop-support.html

upvoted 2 times

Question #259 *Topic 1*

What is the function of BFD?

A. It creates high CPU utilization on hardware deployments

B. It provides uniform failure detection on the same media type

C. It provides uniform failure detection regardless of media type

D. It negotiates to the highest version if the neighbor version differs

Correct Answer: *C*

Community vote distribution

C (100%)

chris110 3 months, 3 weeks ago

Selected Answer: C

The correct answer is:

C. It provides uniform failure detection regardless of media type.

BFD (Bidirectional Forwarding Detection) is a protocol used for rapid failure detection in computer networks. One of its primary functions is to provide uniform and consistent failure detection regardless of the media type or technology being used in the network. BFD can be used with various network media types, including Ethernet, SONET/SDH, MPLS, and more. It ensures that failure detection is fast and consistent, making it a valuable tool for network reliability and fast convergence. Therefore, option C is the correct description of the function of BFD.

upvoted 1 times

Xerath 9 months, 4 weeks ago

Selected Answer: C

https://www.juniper.net/documentation/us/en/software/junos/high-availability/topics/topic-map/bfd.html

BFD can provide fast failure detection times for all media types, encapsulations, topologies, and routing protocols.

upvoted 2 times

Question #260 Topic 1

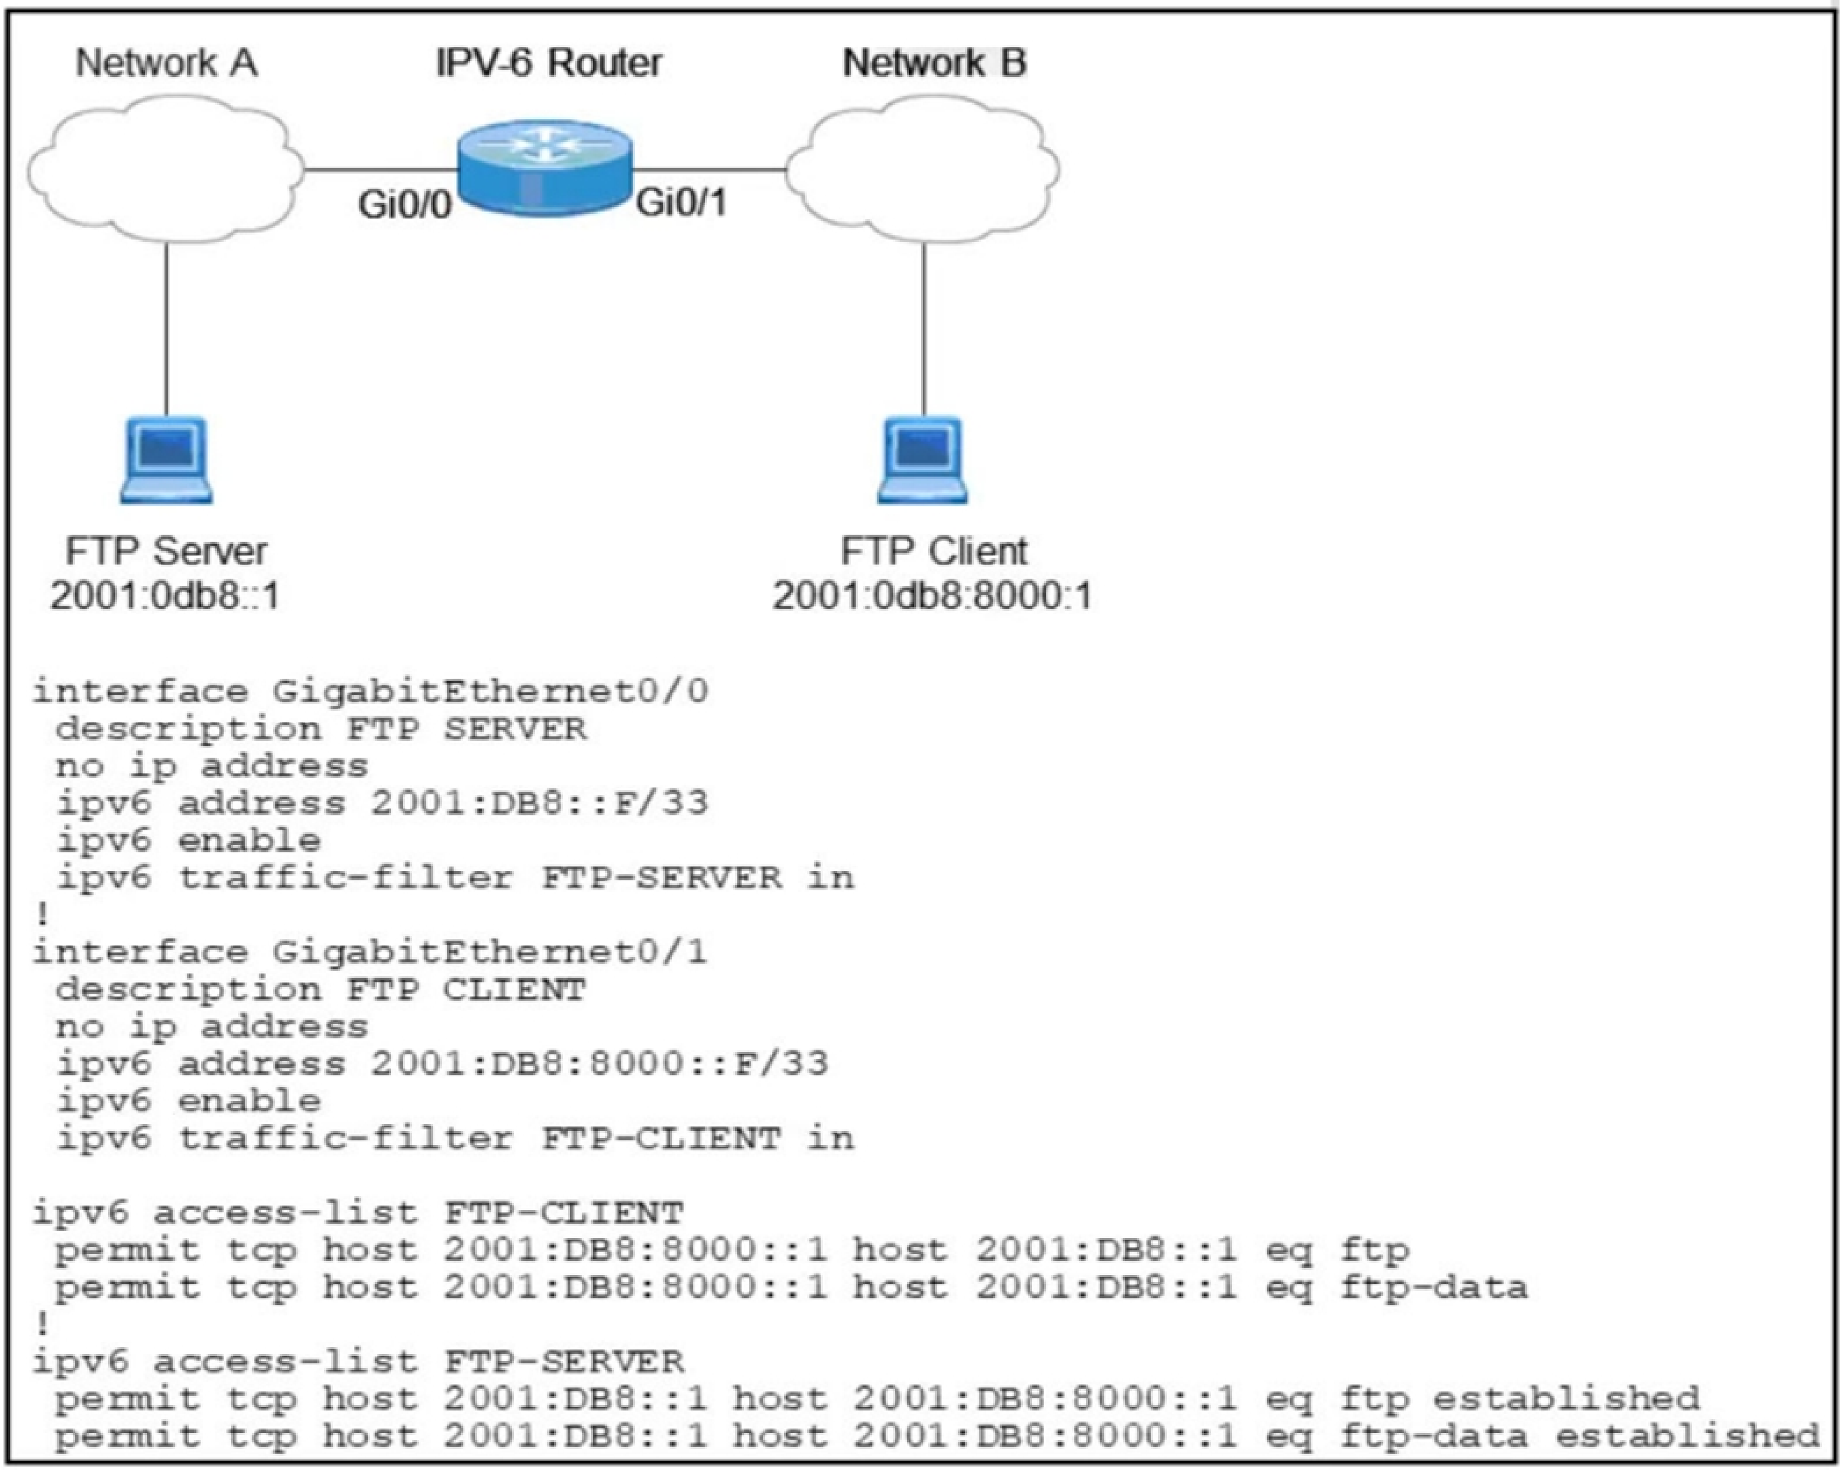

```
interface GigabitEthernet0/0
 description FTP SERVER
 no ip address
 ipv6 address 2001:DB8::F/33
 ipv6 enable
 ipv6 traffic-filter FTP-SERVER in
!
interface GigabitEthernet0/1
 description FTP CLIENT
 no ip address
 ipv6 address 2001:DB8:8000::F/33
 ipv6 enable
 ipv6 traffic-filter FTP-CLIENT in

ipv6 access-list FTP-CLIENT
 permit tcp host 2001:DB8:8000::1 host 2001:DB8::1 eq ftp
 permit tcp host 2001:DB8:8000::1 host 2001:DB8::1 eq ftp-data
!
ipv6 access-list FTP-SERVER
 permit tcp host 2001:DB8::1 host 2001:DB8:8000::1 eq ftp established
 permit tcp host 2001:DB8::1 host 2001:DB8:8000::1 eq ftp-data established
```

Refer to the exhibit. When an FTP client attempts to use passive FTP to connect to the FTP server, the file transfers fail. Which action resolves the issue?

A. Modify traffic filter FTP-SERVER in to the outbound direction.

B. Configure active FTP traffic.

C. Configure to permit TCP ports higher than 1023.

D. Modify FTP-SERVER access list to remove established at the end.

Correct Answer: *C*

Community vote distribution

C (100%)

inteldarvid 5 months, 1 week ago

https://ccnadesdecero.es/diferencias-ftp-modo-activo-pasivo/

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

Option C correct

upvoted 1 times

HungarianDish 7 months, 2 weeks ago

Selected Answer: C

https://community.cisco.com/t5/switching/ftp-and-access-lists/td-p/1525257

upvoted 1 times

DUBC89x 1 year ago

Selected Answer: C

https://serverfault.com/questions/241424/tcp-port-1023-or-tcp-port-1024-for-inbound-passive-ftp-traffic

upvoted 3 times

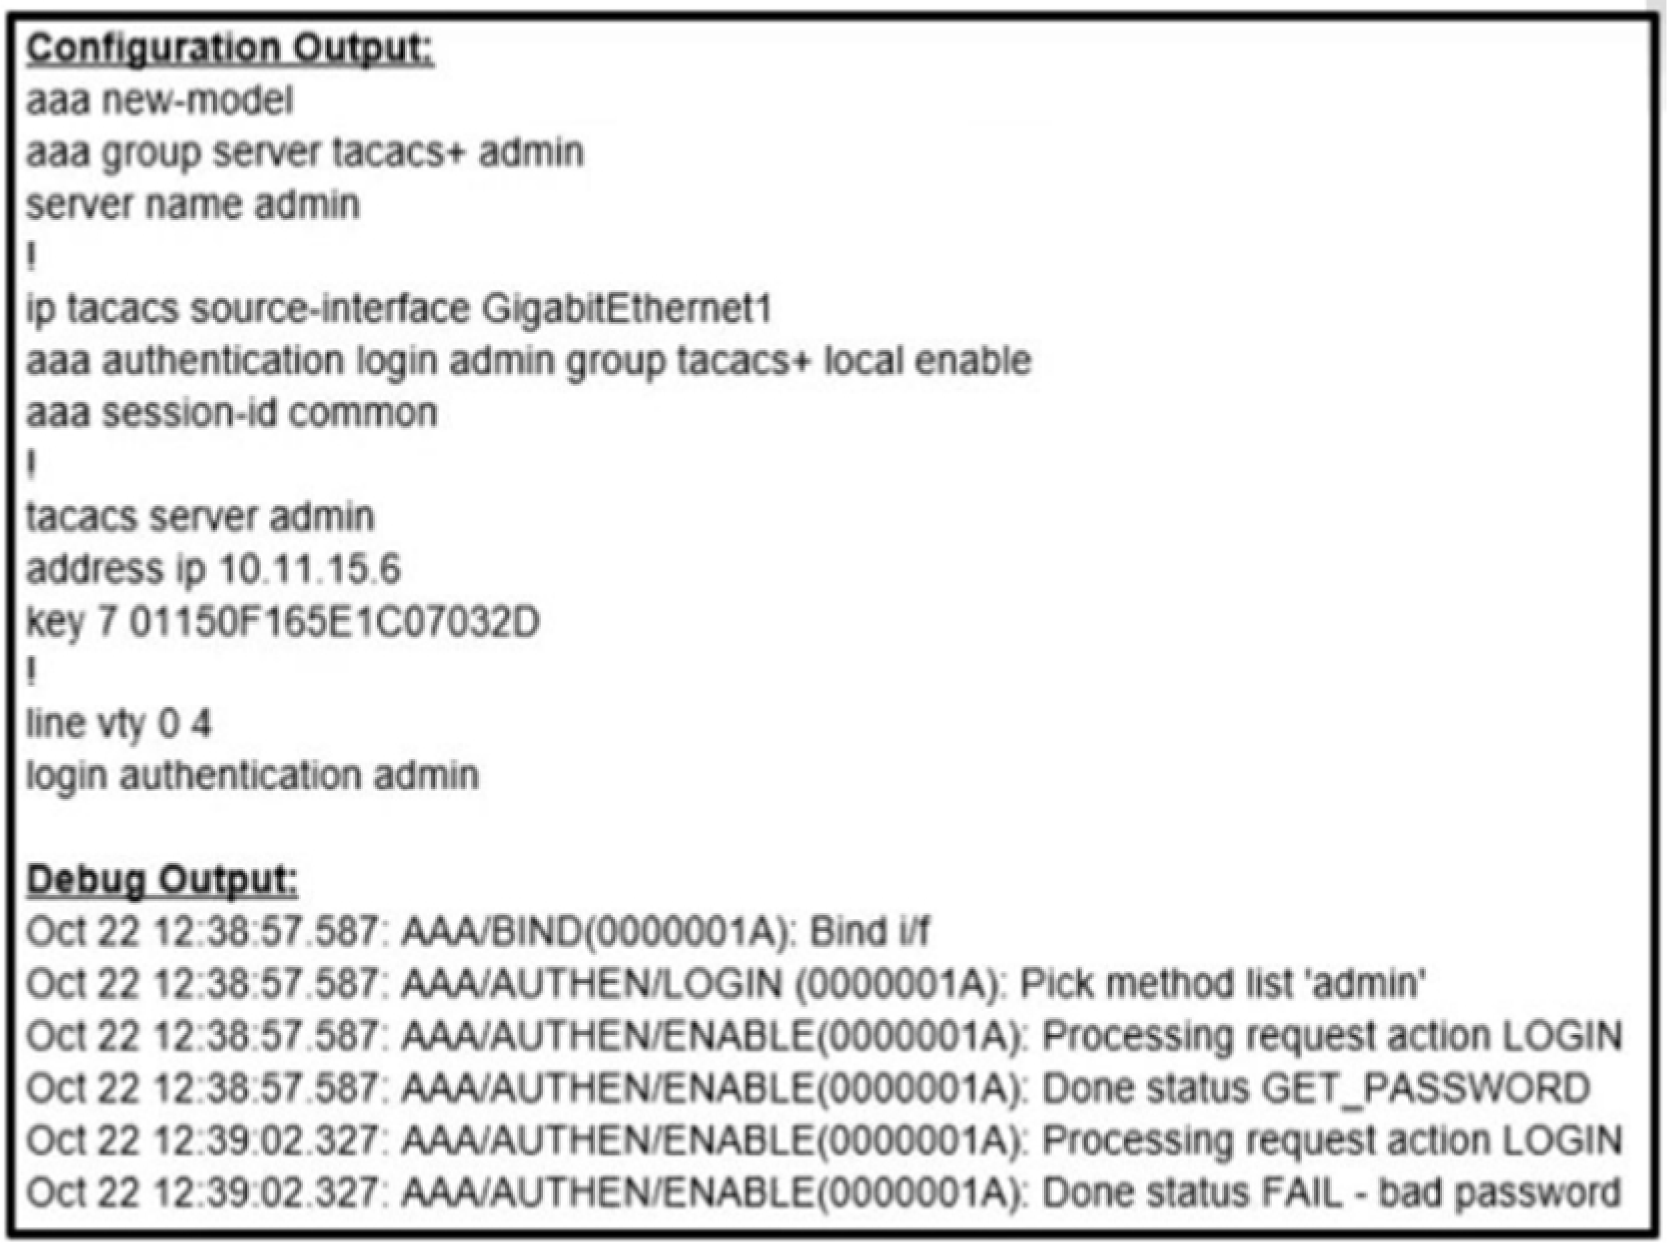

```
Configuration Output:
aaa new-model
aaa group server tacacs+ admin
server name admin
!
ip tacacs source-interface GigabitEthernet1
aaa authentication login admin group tacacs+ local enable
aaa session-id common
!
tacacs server admin
address ip 10.11.15.6
key 7 01150F165E1C07032D
!
line vty 0 4
login authentication admin

Debug Output:
Oct 22 12:38:57.587: AAA/BIND(0000001A): Bind i/f
Oct 22 12:38:57.587: AAA/AUTHEN/LOGIN (0000001A): Pick method list 'admin'
Oct 22 12:38:57.587: AAA/AUTHEN/ENABLE(0000001A): Processing request action LOGIN
Oct 22 12:38:57.587: AAA/AUTHEN/ENABLE(0000001A): Done status GET_PASSWORD
Oct 22 12:39:02.327: AAA/AUTHEN/ENABLE(0000001A): Processing request action LOGIN
Oct 22 12:39:02.327: AAA/AUTHEN/ENABLE(0000001A): Done status FAIL - bad password
```

Refer to the exhibit. An administrator configured a Cisco router for TACACS authentication, but the router is using the local enable password instead. Which action resolves the issue?

A. Configure the aaa authentication login default group admin local if-authenticated command instead.

B. Configure the aaa authentication login admin group tacacs+ local enable none command instead.

C. Configure the aaa authentication login admin group tacacs+ local if-authenticated command instead.

D. Configure the aaa authentication login admin group admin local enable command instead.

Correct Answer: *D*

Reference:

https://community.cisco.com/t5/network-access-control/problem-setting-7606-router-for-tacacs-authentication/td-p/2316903

Community vote distribution

D (86%) 14%

Rob_CCNP000 5 months ago

Selected Answer: D

Correct answer is D the configuration in the exhibit is using a TACACS+ server group called tacacs+ that does not exist. The group is called admin!

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: D

D is correct:

https://community.cisco.com/t5/network-access-control/if-authenticated/td-p/1248124

upvoted 1 times

potato_inet0 7 months, 2 weeks ago

Well, first of all the question seems to be wrong.
We can see the admin method defined and the group is tacacs+ , tacacs server is defined as well as a tacacs server-group.
By applying the aaa authentication login admin group tacacs+ local enable the device should use the defined tacacs server and succesfully communicate, so based on the config there is no issue, I've tested it in LAB.
From the answers D is most logical, the others do not make sense, however the point is the question is wrong.

upvoted 2 times

HungarianDish 7 months, 2 weeks ago

Selected Answer: D

"A" is not reflecting the solution from here:
https://community.cisco.com/t5/network-access-control/problem-setting-7606-router-for-tacacs-authentication/td-p/2316903

"A" adds " if-authenticated", which is used with authorization method lists, and not for authentication.
"D" defines method list "admin" and uses it for "line vty" configuration, which is correct.
Some examples:
https://www.netprojnetworks.com/cisco-9800-tacacs-config-cli-and-verify-notes/

upvoted 2 times

VergilP 1 year, 1 month ago

Selected Answer: D

please review cisco website in jarz 's comment
but I vote for D
the tacacs+ group name is "admin", so it must be "group admin" not "group tacacs+"
so B , C is out
and if-authenticated command is use for aaa authorization
so I choose D

upvoted 1 times

Huntkey 1 year, 2 months ago

Selected Answer: D

I think it is D. The vty line is using the method "admin" and the method "admin" uses the TACACS+ group admin. In the original config, it used a wrong TACACS+ group name that is undefined. Then it doesn't have a local username or password I think. Therefore, causing authentication to refer to the enable password.

upvoted 1 times

Huntkey 1 year, 2 months ago

a little correction. It was using the TACACS+ group "local" and it is undefined. The "local" here is not for using the local credentials

upvoted 1 times

jarz 1 year, 2 months ago

Selected Answer: A

```
aaa authentication login default group admin local enable
```

https://community.cisco.com/t5/network-access-control/problem-setting-7606-router-for-tacacs-authentication/td-p/2316903

upvoted 1 times

VergilP 1 year, 1 month ago

```
aaa authentication login default group admin local enable
```
So You mean answer is D?

upvoted 1 times

VergilP 1 year, 1 month ago

OH , I see the comment below.. in the cisco community

Please replace the below listed command

```
aaa authentication login admin group tacacs+ local enable
```

with;

```
aaa authentication login default group admin local enable
```

upvoted 1 times

Question #262 Topic 1

An administrator attempts to download the .pack NBAR2 file using TFTP from the CPE router to another device over the Gi0/0 interface. The CPE is configured as below: hostname CPE

```
!
ip access-list extended WAN
<`¦>
remark => All UDP rules below for WAN ID: S421T18E58F90
permit udp any eq domain any
permit udp any any eq tftp
deny udp any any
!
interface GigabitEthernet0/0
<`¦>
ip access-group WAN in
<`¦>
!
tftp-server flash:pp-adv-csr1000v-1612.1a-37-53.0.0.pack
```

The transfer fails. Which action resolves this issue?

A. Make the permit udp any eq tftp any entry the last entry in the WAN ACL

B. Shorten the file name to the 8+3 naming convention

C. Change the WAN ACL to permit the entire UDP destination port range

D. Change the WAN ACL to permit the UDP port 69 to allow TFTP

Correct Answer: *C*

Community vote distribution

C (88%) 13%

Huntkey Highly Voted 1 year, 2 months ago

Selected Answer: C

This is actually to my surprise... The TFTP apparently is using the random port for the transfer: TFTP uses UDP as its transport protocol. A transfer request is always initiated targeting port 69, but the data transfer ports are chosen independently by the sender and receiver during the transfer initialization. The ports are chosen at random according to the parameters of the networking stack, typically from the range of ephemeral ports.[4]
https://en.wikipedia.org/wiki/Trivial_File_Transfer_Protocol

upvoted 7 times

[Removed] Most Recent 4 months, 1 week ago

This is interesting. Huntkey provided a nice resource of information, the RFC for TFTP provides explanation as to why this rule actually affects the connection between client and server.
Based on the RFC, TFTP utilizes an ephemeral port named (TID, Transfer Identifier) that is used for the duration of the session. This TID is a random port between 0 to 65535.
When a client sends a Write or Read request (WRQ and RRQ respectively), the Client chooses a TID at random, and sends the request to the server with destination port 69, this is allowed by the ACE #2 in the ACL.
When the server receives the Request, it also chooses a TID at random, and uses that to send the ACK for a WRQ or a the first data packet for RRQ, but this communication is now continued between TIDs as the source/destination UDP ports. this is where the ACE#3 in the ACL is breaking the connection.

1.- CLIENT (src.port.TID) ---(WRQ/RRQ)----> (dst.port.69) TFTP
2.- CLIENT (dst.port.TID) <---(ACK/DATA)--- (src.port.TID) TFTP

upvoted 2 times

inteldarvid 5 months, 1 week ago

Selected Answer: D

D correct:

https://thwack.solarwinds.com/free-tools-trials/f/tftp-server/4613/tftp-communicating-on-high-ports

upvoted 1 times

mrnipsnips 1 year, 1 month ago

This doesn't make sense the ACL is applied 'in' what does it have to do with outbound traffic ?

upvoted 1 times

Question #263

Topic 1

A network administrator must optimize the segment size of the TCP packet on the DMVPN IPsec protected tunnel interface, which carries application traffic from the head office to a designated branch. The TCP segment size must not overwhelm the MTU of the outbound link. Which configuration must be applied to the router to improve the application performance?

A. interface tunnel30 ip mtu 1400 ip tcp payload-size 1360 ! crypto ipsec fragmentation before-encryption

B. interface tunnel30 ip mtu 1400 ip tcp adjust-mss 1360 ! crypto ipsec fragmentation after-encryption

C. interface tunnel30 ip mtu 1400 ip tcp max-segment 1360 ! crypto ipsec fragmentation before-encryption

D. interface tunnel30 ip mtu 1400 ip tcp packet-size 1360 ! crypto ipsec fragmentation after-encryption

Correct Answer: *B*

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_dmvpn/configuration/15-mt/sec-conn-dmvpn-15-mt-book/sec-conn-dmvpn-dmvpn.html

Community vote distribution

B (100%)

HungarianDish 7 months, 2 weeks ago

Selected Answer: B

As well as I see, only "B" contains valid commands.
https://www.networkworld.com/article/2224654/mtu-size-issues.html
https://networkengineering.stackexchange.com/questions/11283/pre-fragmentation-for-ipsec-vpns-on-cisco-routers

upvoted 3 times

Question #264 *Topic 1*

In a DMVPN network, the Spoke1 user observed that the voice traffic is coming to Spoke2 users via the hub router. Which command is required on both spoke routers to communicate directly to one another?

A. ip nhrp nhs multicast

B. ip nhrp shortcut

C. ip nhrp map dynamic

D. ip nhrp redirect

Correct Answer: *B*

Community vote distribution

B (89%) 11%

inteldarvid 5 months, 1 week ago

Selected Answer: B

option B

upvoted 2 times

inteldarvid 5 months, 1 week ago

10000000000000000%%%%% option "B"

upvoted 2 times

HungarianDish 7 months, 2 weeks ago

Selected Answer: B

As well as I see, it is about DMVPN Phase 3 Spoke-to-Spoke Implementation.
Short explanation:
https://carpe-dmvpn.com/2019/02/10/shortcut-dmvpn-demystified/
Long explanation:
https://learningnetwork.cisco.com/s/question/0D53i00000Kt0xkCAB/ip-nhrp-map-multicast-dynamic
https://www.linkedin.com/pulse/dmvpn-i-wish-had-learned-way-from-beginning-leandro-brito
Examples:
https://networkdirection.net/articles/routingandswitching/dmvpn/dmvpn-configuration/
https://www.pearsonitcertification.com/articles/article.aspx?p=3129283&seqNum=8

upvoted 3 times

VergilP 1 year, 1 month ago

Selected Answer: B

agree with ChillingAgain

upvoted 2 times

ChillingAgain 1 year, 1 month ago

Selected Answer: B

Answer is correct. Question is what config is required on both spokes. So not the config on the hub is requested.

upvoted 1 times

chris7890 1 year, 3 months ago

Selected Answer: A

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_dmvpn/configuration/xe-16/sec-conn-dmvpn-xe-16-book/sec-conn-dmvpn-summ-maps.html

ip nhrp nhs [hub-tunnel-ip-address] nbma [hub-wan--ip] multicast - Configures the hub router as the NHRP next-hop server.

upvoted 1 times

Question #265

Topic 1

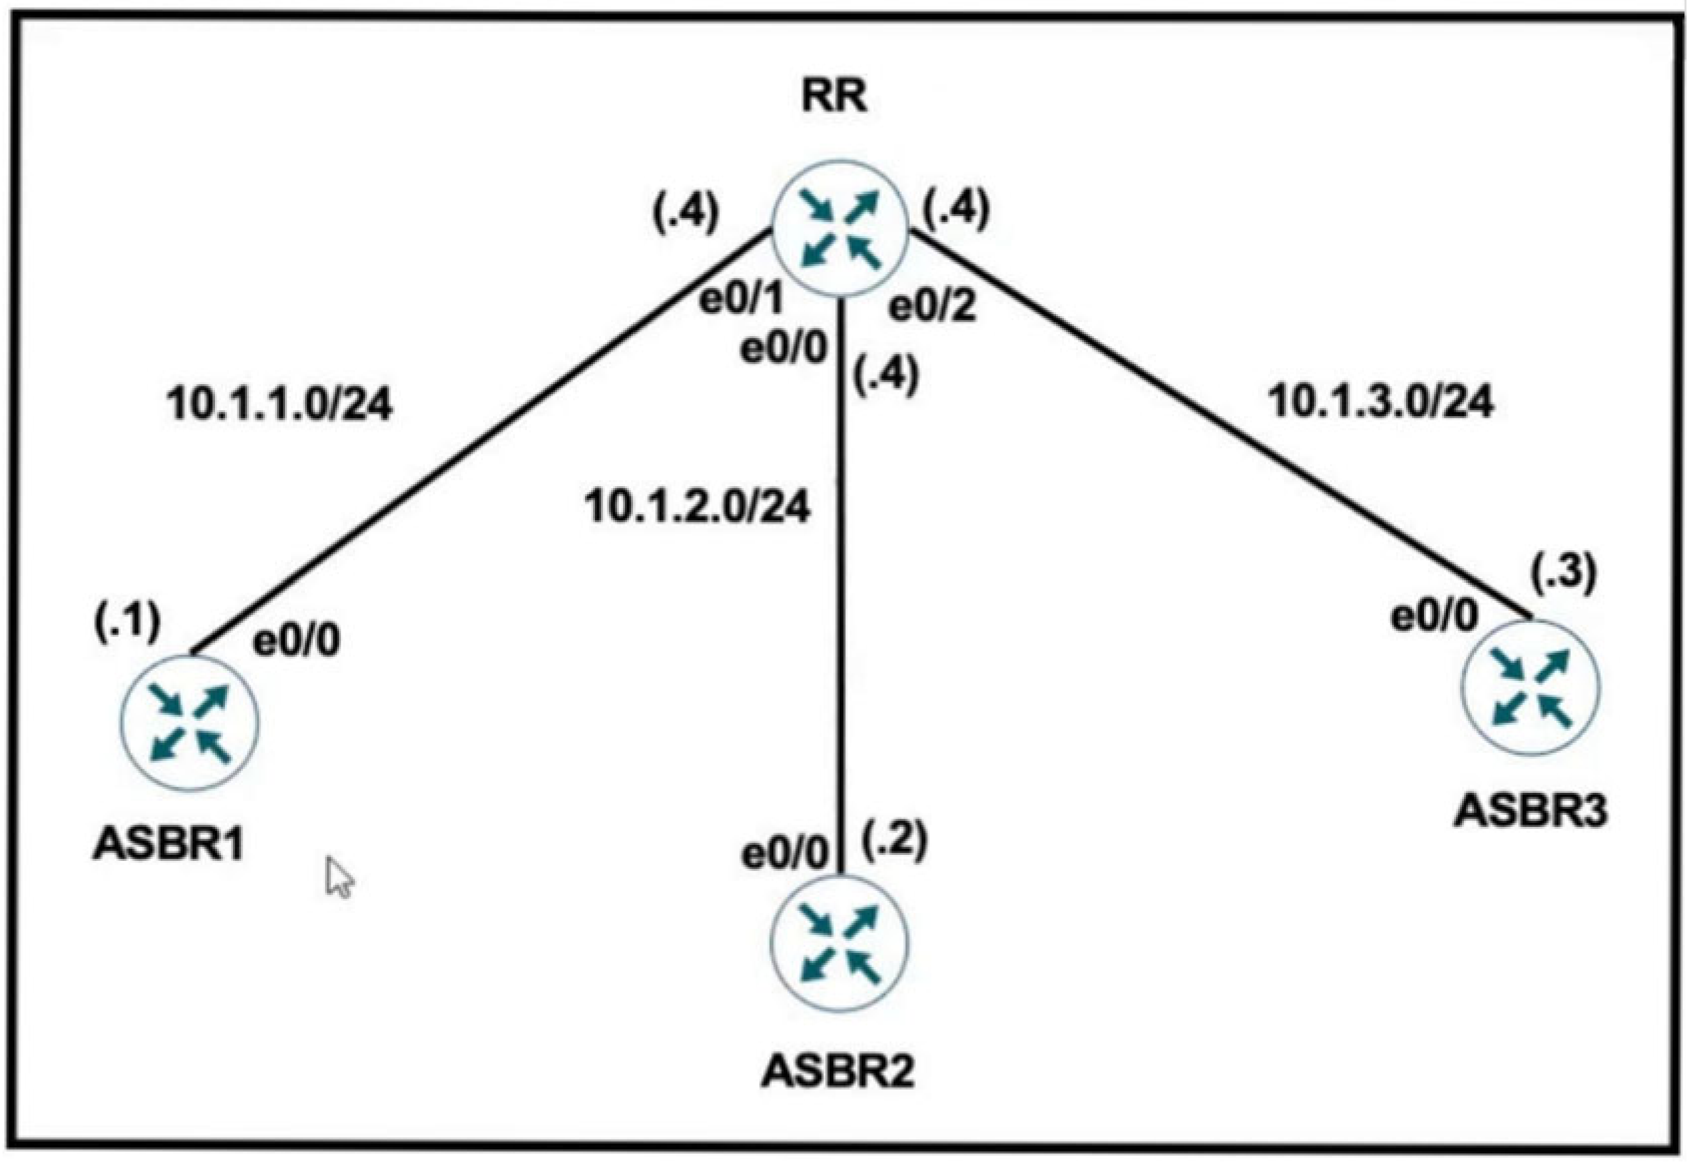

Refer to the exhibit.

RR Configuration:

router bgp 100

neighbor IBGP peer-group

neighbor IBGP route-reflector-client

neighbor 10.1.1.1 remote-as 100

neighbor 10.1.2.2 remote-as 100

neighbor 10.1.3.3 remote-as 100

The network administrator configured the network to establish connectivity between all devices and notices that the ASBRs do not have routes for each other.

Which set of configurations resolves this issue?

A. router bgp 100 neighbor IBGP update-source Loopback0

B. router bgp 100 neighbor IBGP next-hop-self

C. router bgp 100 neighbor 10.1.1.1 next-hop-self neighbor 10.1.2.2 next-hop-self neighbor 10.1.3.3 next-hop-self

D. router bgp 100 neighbor 10.1.1.1 peer-group IBGP neighbor 10.1.2.2 peer-group IBGP neighbor 10.1.3.3 peer-group IBGP

Correct Answer: *D*

Community vote distribution

D (100%)

HungarianDish 7 months, 2 weeks ago

Selected Answer: D

RR is set as the route reflector for the peer-group called IBGP.
For this to take effect, we need to add the neighbors to the perer-group, which is solution "D".
After this, route advertisments will be reflected by RR to the other IBGP routers.

https://community.cisco.com/t5/switching/peer-group-on-a-route-reflector/td-p/1536406
https://networklessons.com/bgp/bgp-route-reflector
https://www.oreilly.com/library/view/cisco-ios-in/0596008694/re638.html

upvoted 4 times

Question #266

Topic 1

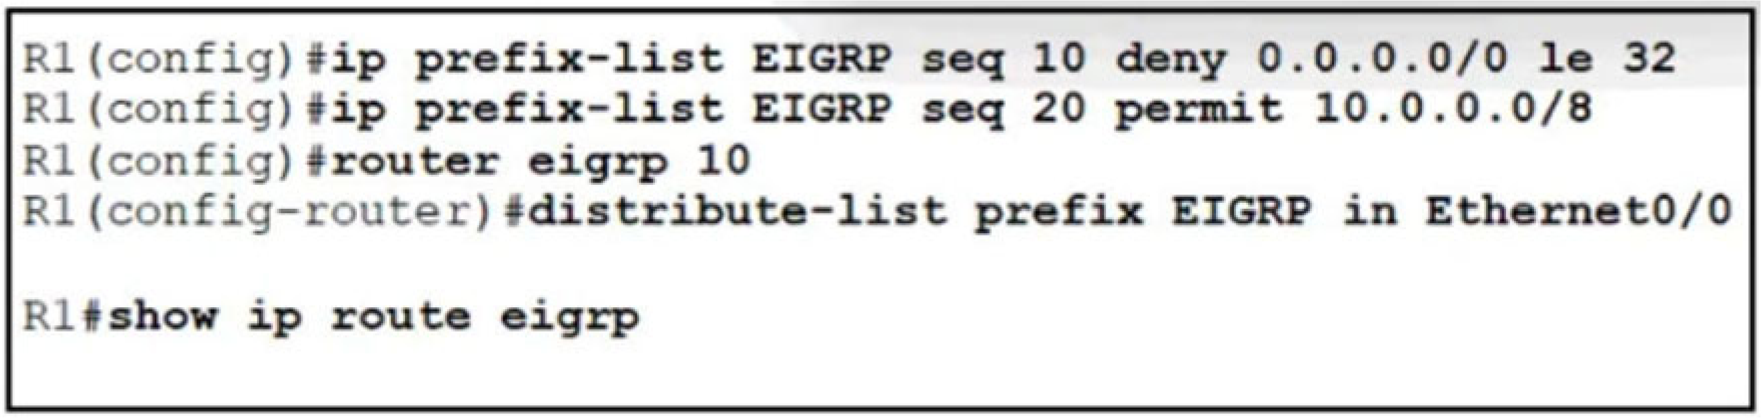

```
R1(config)#ip prefix-list EIGRP seq 10 deny 0.0.0.0/0 le 32
R1(config)#ip prefix-list EIGRP seq 20 permit 10.0.0.0/8
R1(config)#router eigrp 10
R1(config-router)#distribute-list prefix EIGRP in Ethernet0/0

R1#show ip route eigrp
```

Refer to the exhibit. A prefix list is created to filter routes inbound to an EIGRP process except for network 10 prefixes. After the prefix list is applied, no network 10 prefixes are visible in the routing table from EIGRP. Which configuration resolves the issue?

A. ip prefix-list EIGRP seq 10 permit 0.0.0.0/0 le 32

B. ip prefix-list EIGRP seq 20 permit 10.0.0.0/8 ge 9 ip prefix-list EIGRP seq 10 permit 0.0.0.0/0 le 32

C. ip prefix-list EIGRP seq 20 permit 10.0.0.0/8 ge 9

D. ip prefix-list EIGRP seq 5 permit 10.0.0.0/8 ge 9 no ip prefix-list EIGRP seq 20 permit 10.0.0.0/8

Correct Answer: *D*

Community vote distribution

D (100%)

HungarianDish 7 months, 2 weeks ago

Selected Answer: D

"ip prefix-list EIGRP seq 5 permit 10.0.0.0/8" is correct. A prefix-list is an ordered list. "permit 10.0.0.0/8" needs to come before "deny 0.0.0.0/0 le 32" (deny everything), otherwise the "10" network is matched by the deny statement and thus, it gets to be filtered. "sequence 5" places the "permit 10.0.0.0/8" before "deny 0.0.0.0/0 le 32".

upvoted 3 times

[Removed] 4 months, 2 weeks ago

To add to this answer, the second problem resolved in answer D is the acceptance of prefix lengths greater than /8. As it stood, sequence 20 was only accepting the prefix "10.0.0.0/8" and nothing else. The keyword "ge 9" allows the prefix statement to accept prefix lengths between /8 and /32. Alternatively it could have been "le 32"

upvoted 1 times

HungarianDish 7 months, 2 weeks ago

https://networklessons.com/eigrp/how-to-configure-prefix-list-on-cisco-router

upvoted 1 times

Question #267 *Topic 1*

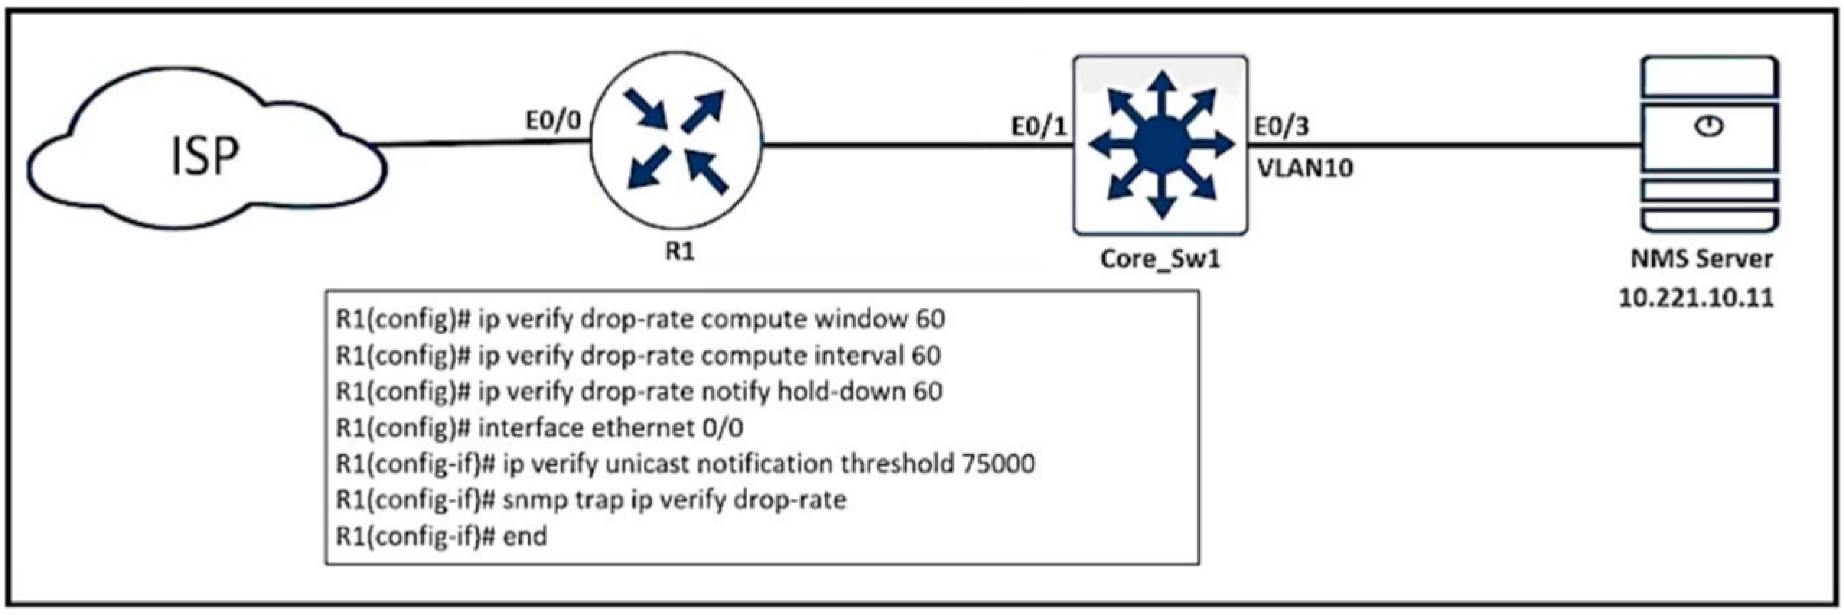

```
R1(config)# ip verify drop-rate compute window 60
R1(config)# ip verify drop-rate compute interval 60
R1(config)# ip verify drop-rate notify hold-down 60
R1(config)# interface ethernet 0/0
R1(config-if)# ip verify unicast notification threshold 75000
R1(config-if)# snmp trap ip verify drop-rate
R1(config-if)# end
```

Refer to the exhibit. An engineer configured SNMP traps to record spoofed packets drop of more than 48000 a minute on the ethernet0/0 interface. During an IP spoofing attack, the engineer noticed that no notifications have been received by the SNMP server. Which configuration resolves the issue on R1?

A. ip verify unicast notification threshold 800

B. ip verify unicast notification threshold 8000

C. ip verify unicast notification threshold 48000

D. ip verify unicast notification threshold 80

Correct Answer: *A*

Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_urpf/configuration/12-4t/sec-data-urpf-12-4t-book/sec-urpf-mib.html

Community vote distribution

A (100%)

Slinky Highly Voted 7 months, 4 weeks ago

Selected Answer: A

The "ip verify unicast notification threshold 800" command specifies the number of packets per second. So in this case, 800 packets a second X 60 seconds in a minute, you get 48,000 packets.

upvoted 11 times

HungarianDish 7 months, 2 weeks ago

Great explanation!

upvoted 2 times

Jey117 Most Recent 2 months, 1 week ago

How the hello are we supposed to know this sh1t? We don't work at Cisco TAC

upvoted 2 times

Juniour 10 months, 2 weeks ago

correct

upvoted 1 times

Question #268 *Topic 1*

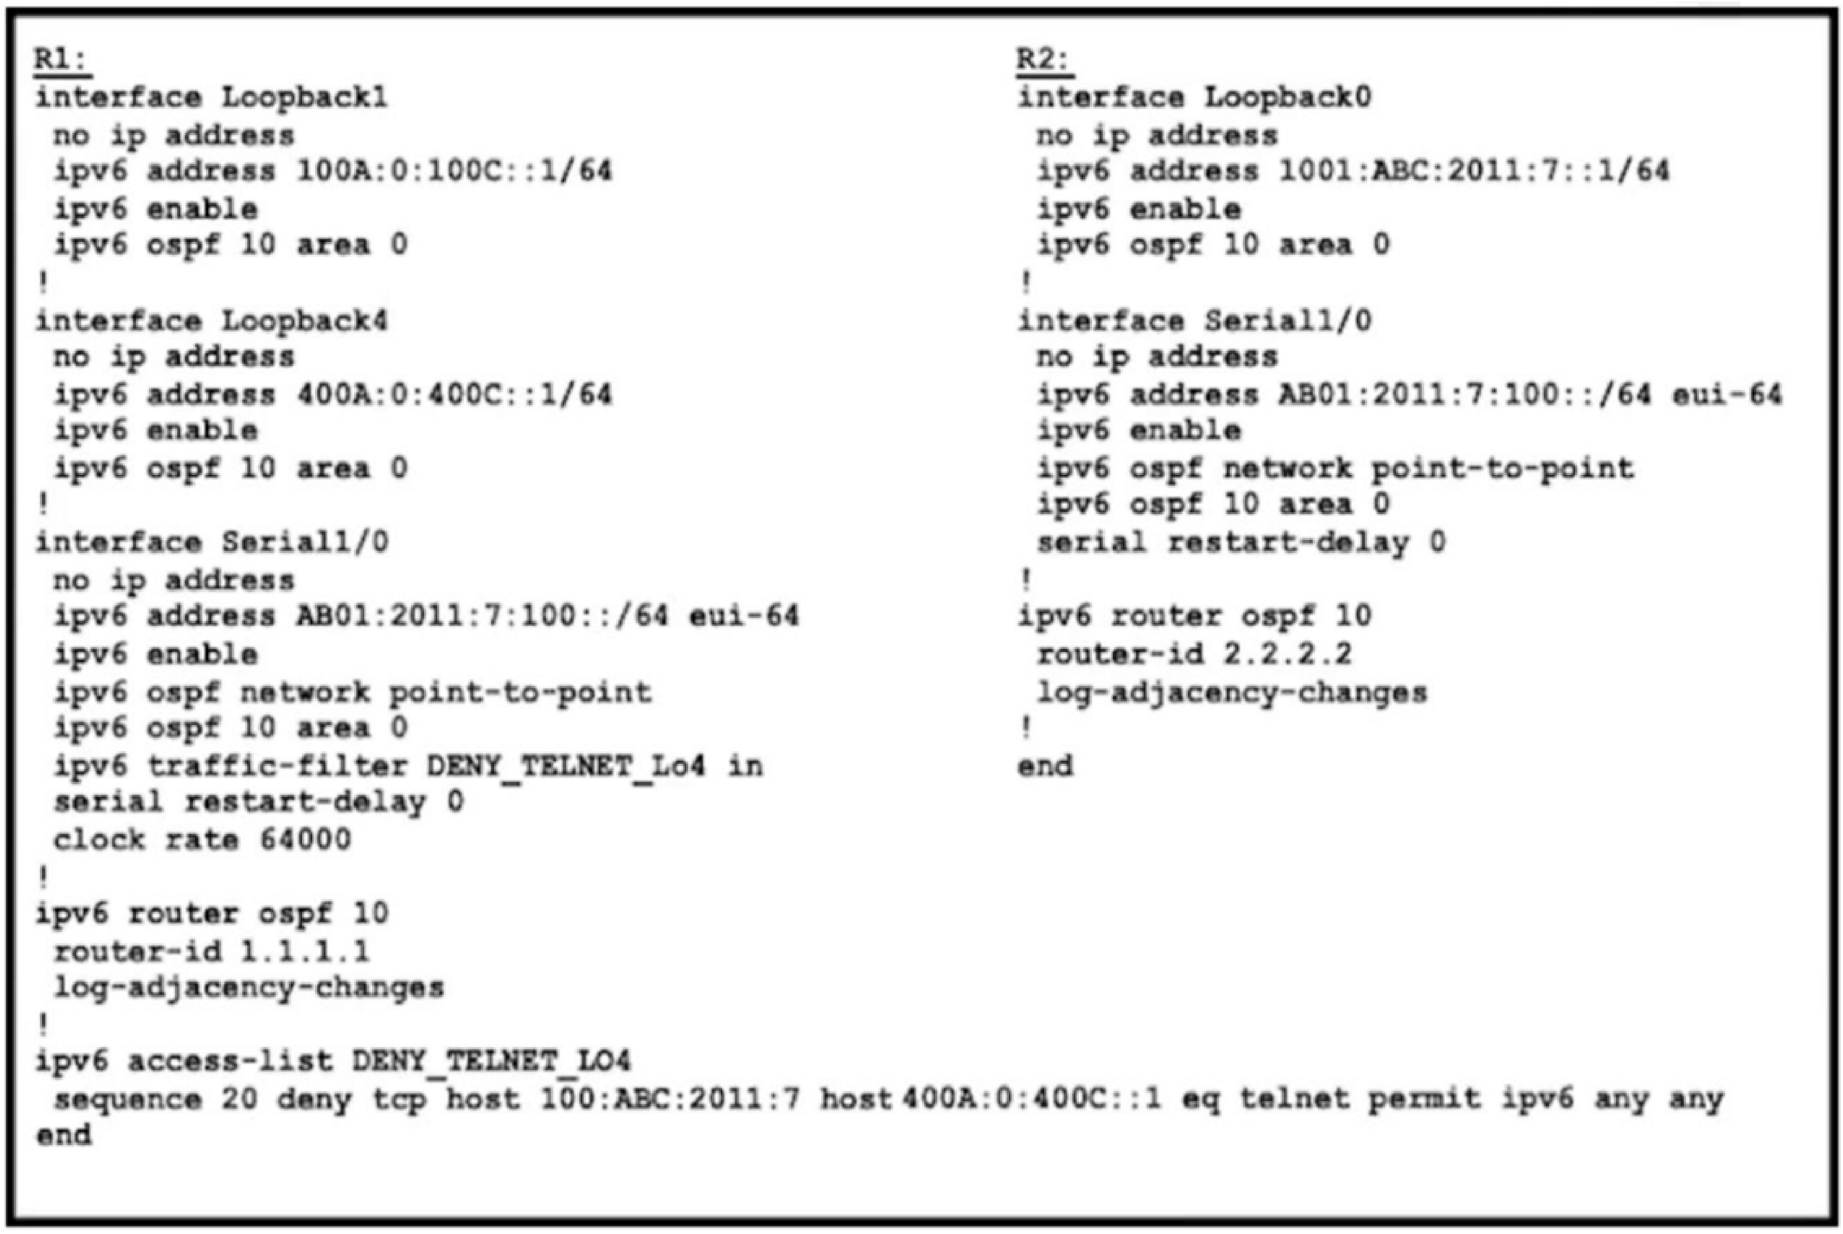

```
R1:
interface Loopback1
 no ip address
 ipv6 address 100A:0:100C::1/64
 ipv6 enable
 ipv6 ospf 10 area 0
!
interface Loopback4
 no ip address
 ipv6 address 400A:0:400C::1/64
 ipv6 enable
 ipv6 ospf 10 area 0
!
interface Serial1/0
 no ip address
 ipv6 address AB01:2011:7:100::/64 eui-64
 ipv6 enable
 ipv6 ospf network point-to-point
 ipv6 ospf 10 area 0
 ipv6 traffic-filter DENY_TELNET_Lo4 in
 serial restart-delay 0
 clock rate 64000
!
ipv6 router ospf 10
 router-id 1.1.1.1
 log-adjacency-changes
!
ipv6 access-list DENY_TELNET_LO4
 sequence 20 deny tcp host 100:ABC:2011:7 host 400A:0:400C::1 eq telnet permit ipv6 any any
end
```

```
R2:
interface Loopback0
 no ip address
 ipv6 address 1001:ABC:2011:7::1/64
 ipv6 enable
 ipv6 ospf 10 area 0
!
interface Serial1/0
 no ip address
 ipv6 address AB01:2011:7:100::/64 eui-64
 ipv6 enable
 ipv6 ospf network point-to-point
 ipv6 ospf 10 area 0
 serial restart-delay 0
!
ipv6 router ospf 10
 router-id 2.2.2.2
 log-adjacency-changes
!
end
```

Refer to the exhibit. An engineer implemented an access list on R1 to allow anyone to Telnet except R2 Loopback0 to R1 Loopback4. How must sequence 20 be replaced on the R1 access list to resolve the issue?

A. sequence 20 permit tcp host 1001:ABC:2011:7::1 host 400A:0:400C::1 eq telnet

B. sequence 20 deny tcp host 400A:0:400C::1 host 1001:ABC:2011:7::1 eq telnet

C. sequence 20 permit tcp host 400A:0:400C::1 host 1001:ABC:2011:7::1 eq telnet

D. sequence 20 deny tcp host 1001:ABC:2011:7::1 host 400A:0:400C::1 eq telnet

Correct Answer: *D*

Community vote distribution

D (100%)

HungarianDish 7 months, 2 weeks ago

Selected Answer: D

Agree with solution "D". source ipv6 address needs to be corrected to "1001:ABC:2011:7::1". The "deny" statement is required.

upvoted 3 times

Question #269 *Topic 1*

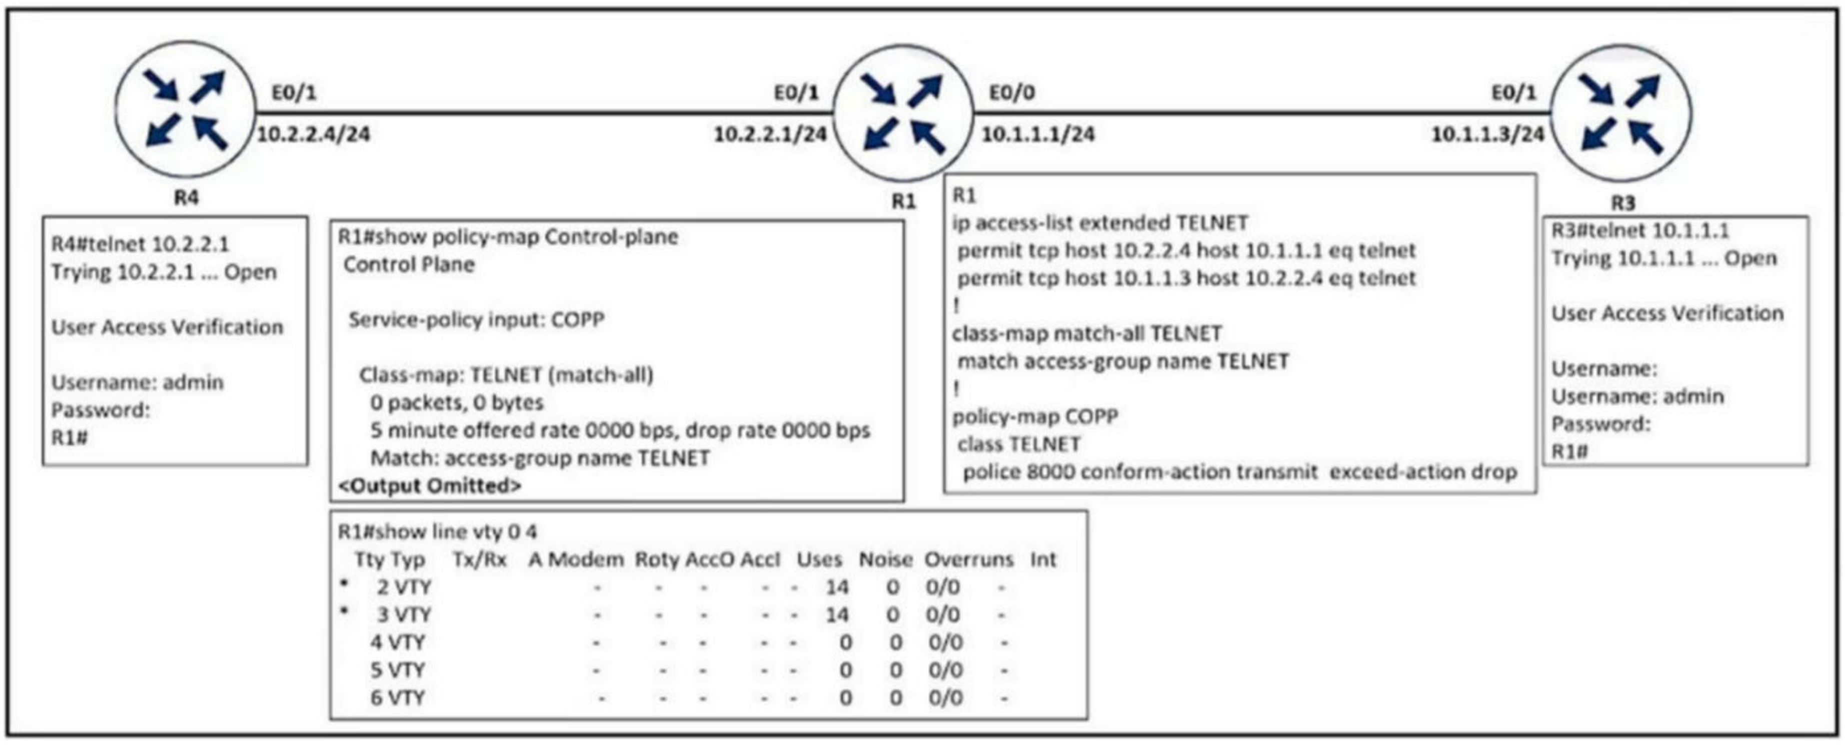

```
R4#telnet 10.2.2.1
Trying 10.2.2.1 ... Open

User Access Verification

Username: admin
Password:
R1#
```

```
R1#show policy-map Control-plane
 Control Plane

  Service-policy input: COPP

   Class-map: TELNET (match-all)
    0 packets, 0 bytes
    5 minute offered rate 0000 bps, drop rate 0000 bps
    Match: access-group name TELNET
<Output Omitted>
```

```
R1
ip access-list extended TELNET
 permit tcp host 10.2.2.4 host 10.1.1.1 eq telnet
 permit tcp host 10.1.1.3 host 10.2.2.4 eq telnet
!
class-map match-all TELNET
 match access-group name TELNET
!
policy-map COPP
 class TELNET
  police 8000 conform-action transmit  exceed-action drop
```

```
R3#telnet 10.1.1.1
Trying 10.1.1.1 ... Open

User Access Verification

Username:
Username: admin
Password:
R1#
```

```
R1#show line vty 0 4
  Tty Typ    Tx/Rx    A Modem  Roty AccO Accl   Uses   Noise  Overruns   Int
*    2 VTY                  -       -     -        -  -    14      0     0/0      -
*    3 VTY                  -       -     -        -  -    14      0     0/0      -
     4 VTY                  -       -     -        -  -     0      0     0/0      -
     5 VTY                  -       -     -        -  -     0      0     0/0      -
     6 VTY                  -       -     -        -  -     0      0     0/0      -
```

Refer to the exhibit. An engineer implemented CoPP to limit Telnet traffic to protect the router CPU. It was noticed that the Telnet traffic did not pass through

CoPP. Which configuration resolves the issue?

A. ip access-list extended TELNET permit tcp host 10.2.2.1 host 10.2.2.4 eq telnet permit tcp host 10.1.1.1 host 10.1.1.3 eq telnet

B. policy-map COPP class TELNET police 8000 conform-action transmit exceed-action transmit

C. ip access-list extended TELNET permit tcp host 10.2.2.4 host 10.2.2.1 eq telnet permit tcp host 10.1.1.3 host 10.1.1.1 eq telnet

D. policy-map COPP class TELNET police 8000 conform-action transmit exceed-action transmit violate-action drop

Correct Answer: *C*

Community vote distribution

C (75%)

D (25%)

guy276465281819372 5 months ago

Selected Answer: C

C is correct. matching IP address source and destination.

upvoted 1 times

HungarianDish 7 months, 2 weeks ago

I meant the destination IPs in the access-list. Destination IPs need to be corrected.

upvoted 1 times

HungarianDish 7 months, 2 weeks ago

Selected Answer: C

"exceed-action drop" achieves the goal, however, the source IPs in the access-list are wrong and need to be corrected for sure. So it is "C" for me.

upvoted 1 times

GodFather 10 months, 3 weeks ago

Selected Answer: C

police bps [burst-normal] [burst-max] conform-action action exceed-action action [violate-action action]
Syntax Description

bps
Average rate, in bits per second. Valid values are 8000 to 200000000.
burst-normal
(Optional) Normal burst size in bytes. Valid values are 1000 to 51200000. Default normal burst size is 1500.
burst-max
(Optional) Maximum burst size, in bytes. Valid values are 1000 to 51200000. Default varies by platform.
conform-action
Specifies action to take on packets that conform to the rate limit.
exceed-action
Specifies action to take on packets that exceed the rate limit.
violate-action
(Optional) Specifies action to take on packets that violate the normal and maximum burst sizes.
action
Action to take on packets. Specify one of the following keywords:
•drop—Drops the packet.

upvoted 1 times

herojacky 11 months, 3 weeks ago

Selected Answer: D

limit Telnet traffic to protect the router CPU

upvoted 1 times

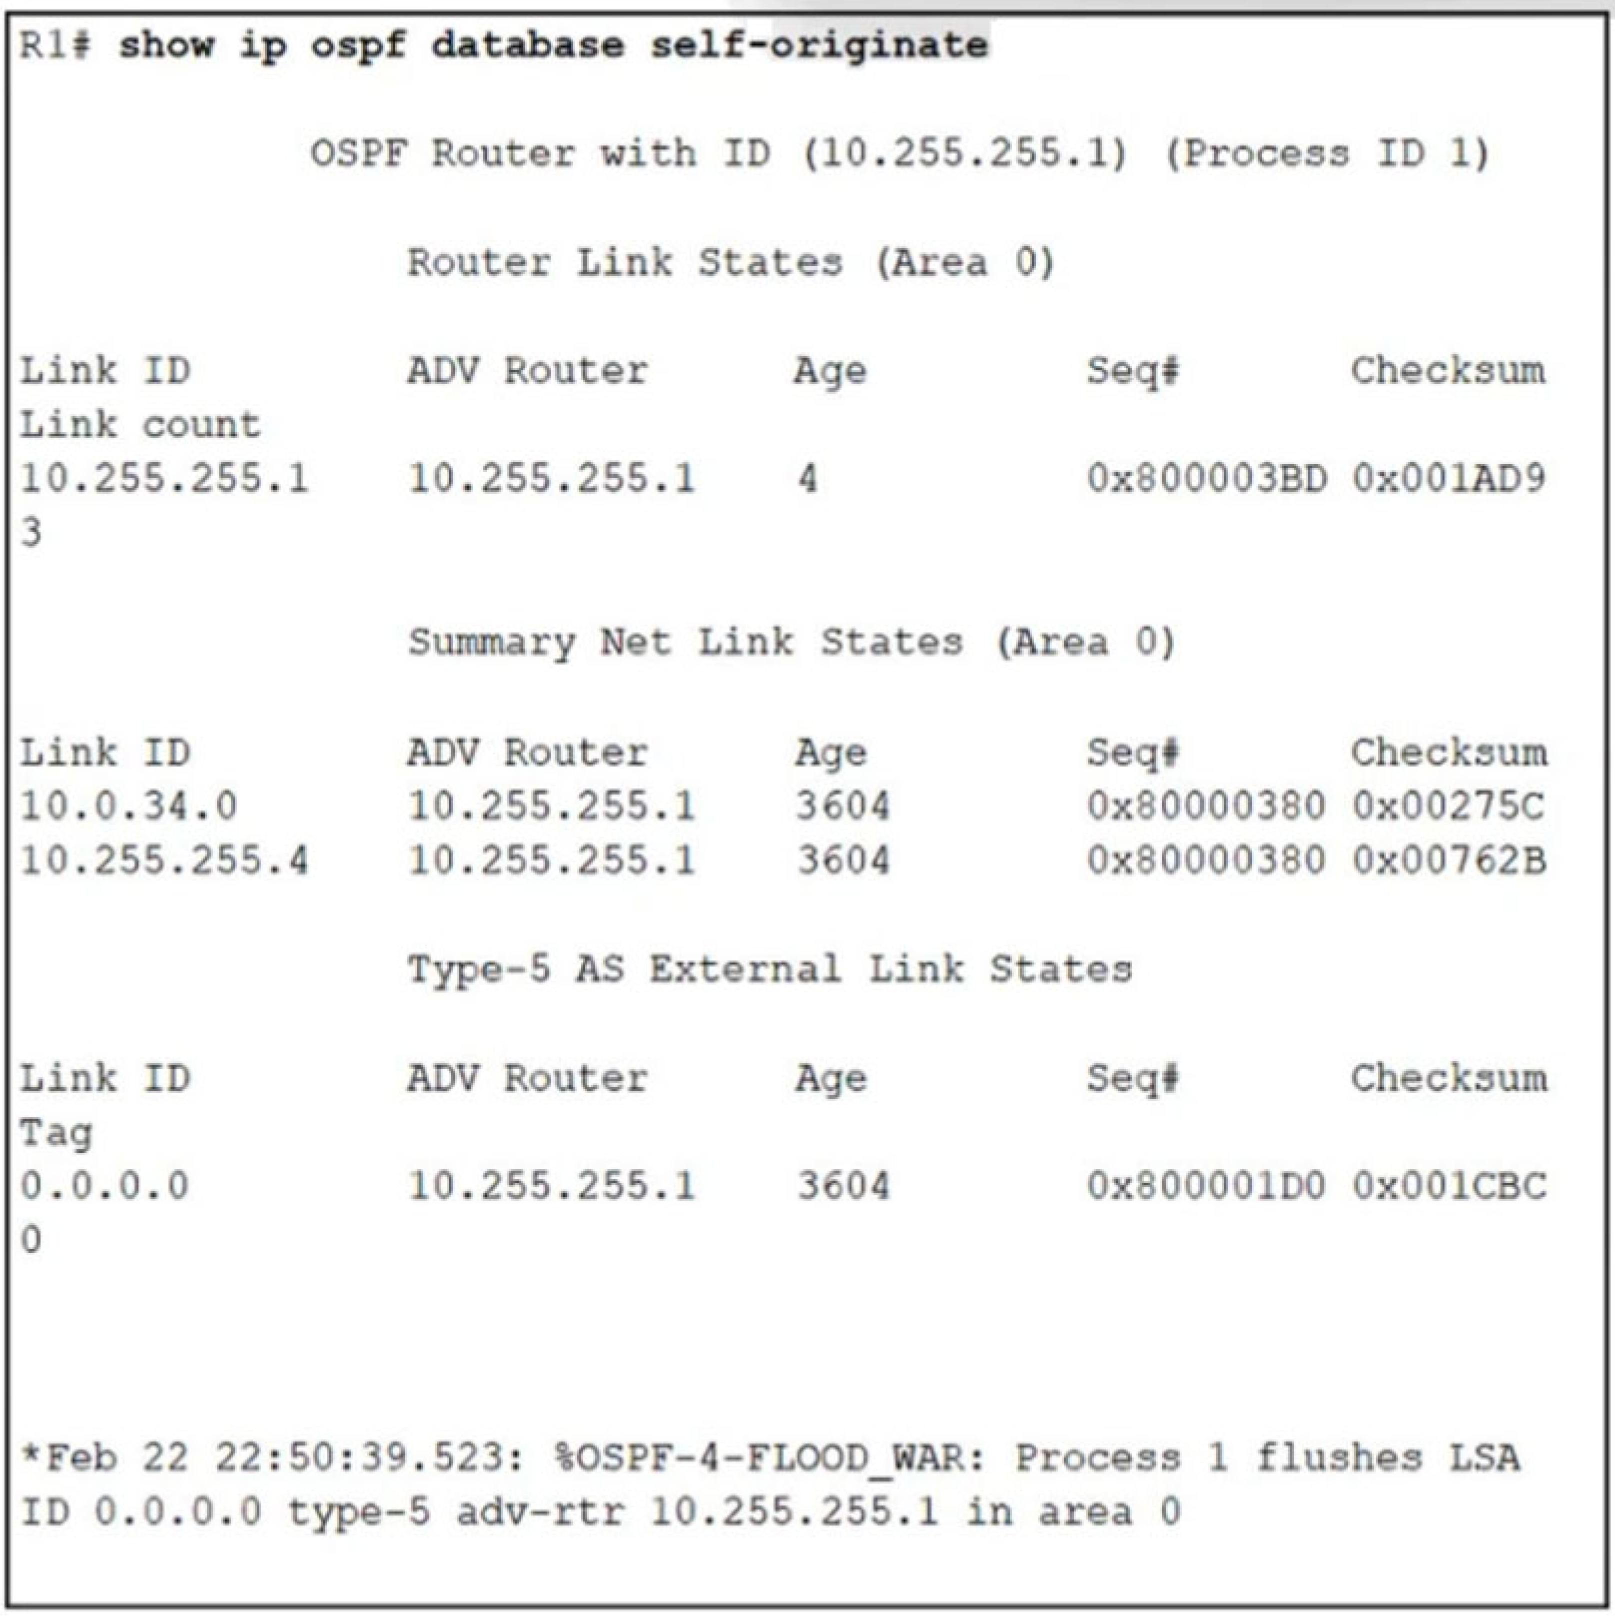

```
R1# show ip ospf database self-originate

            OSPF Router with ID (10.255.255.1) (Process ID 1)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
Link count
10.255.255.1    10.255.255.1    4           0x800003BD 0x001AD9
3

                Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.0.34.0       10.255.255.1    3604        0x80000380 0x00275C
10.255.255.4    10.255.255.1    3604        0x80000380 0x00762B

                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum
Tag
0.0.0.0         10.255.255.1    3604        0x800001D0 0x001CBC
0



*Feb 22 22:50:39.523: %OSPF-4-FLOOD_WAR: Process 1 flushes LSA
ID 0.0.0.0 type-5 adv-rtr 10.255.255.1 in area 0
```

Refer to the exhibit. After configuring OSPF in R1, some external destinations in the network became unreachable. Which action resolves the issue?

A. Disconnect the router with the OSPF router ID 0.0.0 0 from the network.

B. Increase the SPF delay interval on R1 to synchronize routes.

C. Change the R1 router ID from 10.255.255.1 to a unique value and clear the process.

D. Clear the OSPF process on R1 to flush stale LSAs sent by other routers.

Correct Answer: *C*

Reference:

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/118880-technote-ospf-00.html

Community vote distribution

C (100%)

HungarianDish 7 months, 2 weeks ago

Selected Answer: C

The OSPF Router ID 10.255.255.1 is not unique, thus "OSPF-4-FLOOD_WAR" error message is generated on the affected routers.
R1 is one of the affected devices, so "C) Change the R1 router ID from 10.255.255.1 to a unique value and clear the process" resolves the issue.

upvoted 4 times

HungarianDish 7 months, 2 weeks ago

"show ip ospf database self-originate" displays LSAs from the local router = R1.
R1 has the OSPF Router ID 10.255.255.1 (displayed as "adv-rtr" or "ADV Router").

As we see, R1 originates a type 5 LSA with a link ID of 0.0.0.0, which is the default route (from default-information originate).
That is where the router ID conflict occurs.
upvoted 2 times

HungarianDish 7 months, 2 weeks ago

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/118880-technote-ospf-00.html
OSPF-4-FLOOD_WAR
"...Type-5 LSAs when there is a duplicate router ID in different OSPF Areas"
upvoted 2 times

HungarianDish 7 months, 2 weeks ago

https://community.cisco.com/t5/switching/ospf-4-flood-war-messages-after-config-change/td-p/2506500
"For OSPF to function correctly the IP addresses of transit networks must be unique.
If it is not unique the conflicting routers reports this error message.
In the error message the router with the OSPF router ID reported as adv-rtr reports this message."
upvoted 2 times

bolbolskanes 12 months ago

C correct answer
Ref: https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/9237-9.html
upvoted 2 times

Question #271 Topic 1

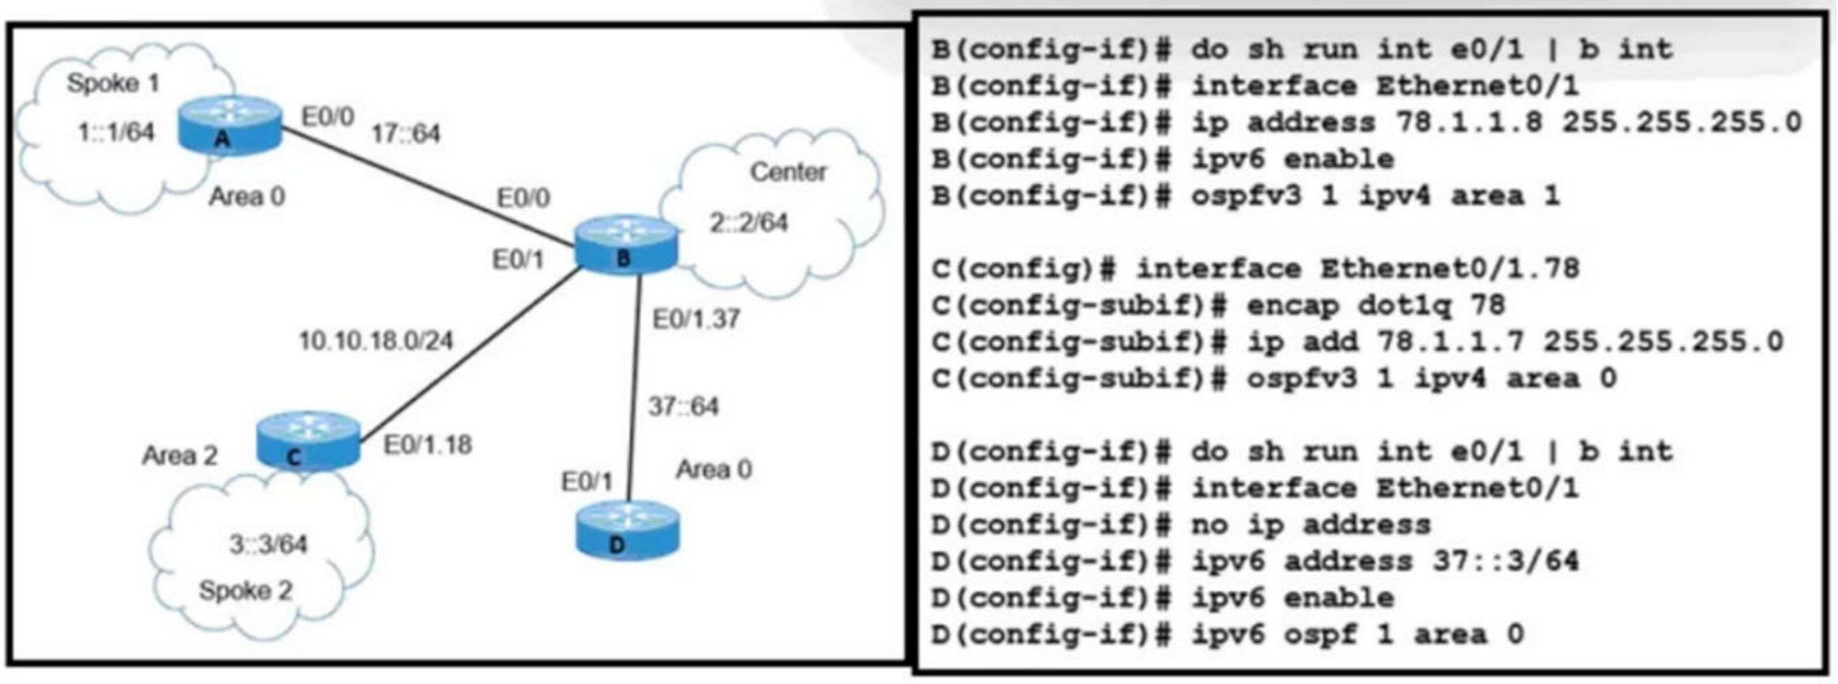

```
B(config-if)# do sh run int e0/1 | b int
B(config-if)# interface Ethernet0/1
B(config-if)# ip address 78.1.1.8 255.255.255.0
B(config-if)# ipv6 enable
B(config-if)# ospfv3 1 ipv4 area 1

C(config)# interface Ethernet0/1.78
C(config-subif)# encap dot1q 78
C(config-subif)# ip add 78.1.1.7 255.255.255.0
C(config-subif)# ospfv3 1 ipv4 area 0

D(config-if)# do sh run int e0/1 | b int
D(config-if)# interface Ethernet0/1
D(config-if)# no ip address
D(config-if)# ipv6 address 37::3/64
D(config-if)# ipv6 enable
D(config-if)# ipv6 ospf 1 area 0
```

Refer to the exhibit. A network engineer receives a report that Spoke 1 users can perform bank transactions with the server located at the Center site, but Spoke 2 users cannot. Which action resolves the issue?

A. Configure the Spoke 2 users IP on the router B OSPF domain

B. Configure IPv6 on the routers B and C interfaces

C. Configure OSPFv2 on the routers B and C interfaces

D. Configure encapsulation dot1q 78 on the router C interface

Correct Answer: *B*

Community vote distribution

B (100%)

⊟ **fizzer** 3 months, 1 week ago

I agree that this question is somewhat stupid, the interface will not even take the ospfv3 configuration command unless ipv6 is already enabled on the interface, not sure how they managed to display and interface config with the ospfv3 command already in place without the ipv6 enable

```
ABR0-1(config-if)#ospfv3 1 ipv4 area 0
% OSPFv3: IPV6 is not enabled on this interface
ABR0-1(config-if)#ipv6 enable
ABR0-1(config-if)#ospfv3 1 ipv4 area 0
ABR0-1(config-if)#^Z
```

upvoted 2 times

⊟ **[Removed]** 4 months, 4 weeks ago

Fucking idiotic question. It tests nothing in terms of knowledge. How do you apply for the job of making questions for Cisco exams, seems like an easy job...

upvoted 4 times

⊟ **guy276465281819372** 5 months ago

That is the most bizarre and stupid question I have ever read.

upvoted 1 times

⊟ **Rob_CCNP000** 6 months ago

None of these answers would really fix the problem. Who writes these question! Absolutely terrible.

upvoted 2 times

⊟ **Almylle** 5 months, 4 weeks ago

If it's this dumps are really valid, im really dissapointed with cisco, like the 90% of the questions at this question are horrible

upvoted 1 times

⊟ **HungarianDish** 6 months, 2 weeks ago

Selected Answer: B

Answer "B". abd123 is right, "ipv6 enable" is missing for ospfv3.

upvoted 1 times

⊟ **HungarianDish** 7 months, 2 weeks ago

My guesses:

The connecting routers should have one leg in OSPF area 0. Certainly, the interfaces for connection B-C should be in OSPF area 0, and that is

missing on Router B.
All interfaces in IPv6 OSPFv3 domain should have an IPv6 address, and that is missing on Router C.
For me it looks like a frame relay topology, so probably the encapsulation should be frame relay.
upvoted 1 times

abd123 10 months, 2 weeks ago

Selected Answer: B

using ospf v3 you need IPV6 enable
upvoted 4 times

dq28 11 months, 3 weeks ago

So many problems to see here! Area-Mismatch, maybe an encapsulation mismatch and yes also IPv6 is also an problem here. But none of the answers make sense in this case!
upvoted 1 times

VergilP 1 year, 1 month ago

Agree with ChillingAgain
can someone explain this question?
upvoted 1 times

ChillingAgain 1 year, 1 month ago

Badly written question? Cannot understand what would be a valid option. Any ideas, someone?
upvoted 2 times

Question #272 *Topic 1*

What is an MPLS LDP targeted session?

A. LDP session established by exchanging multicast hello packets

B. LDP session established between LSRs by exchanging TCP hello packets

C. session between neighbors that are connected no more than one hop away

D. label distribution session between non-directly connected neighbors

Correct Answer: *D*

Community vote distribution

D (100%)

Ckl22 1 year ago

Selected Answer: D

Given answer is correct.

https://community.cisco.com/t5/mpls/targeted-ldp-sessions/td-p/2288569
upvoted 3 times

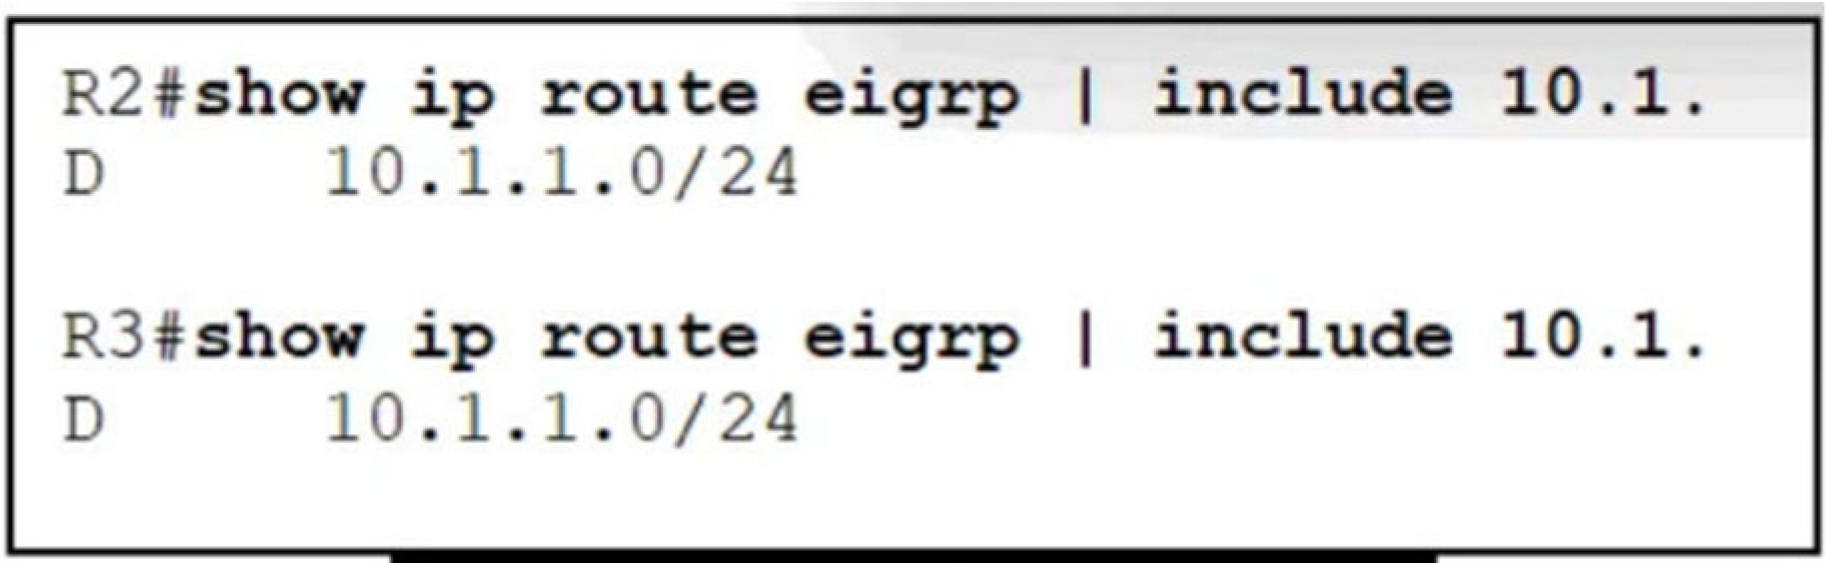

```
R2#show ip route eigrp | include 10.1.
D      10.1.1.0/24

R3#show ip route eigrp | include 10.1.
D      10.1.1.0/24
```

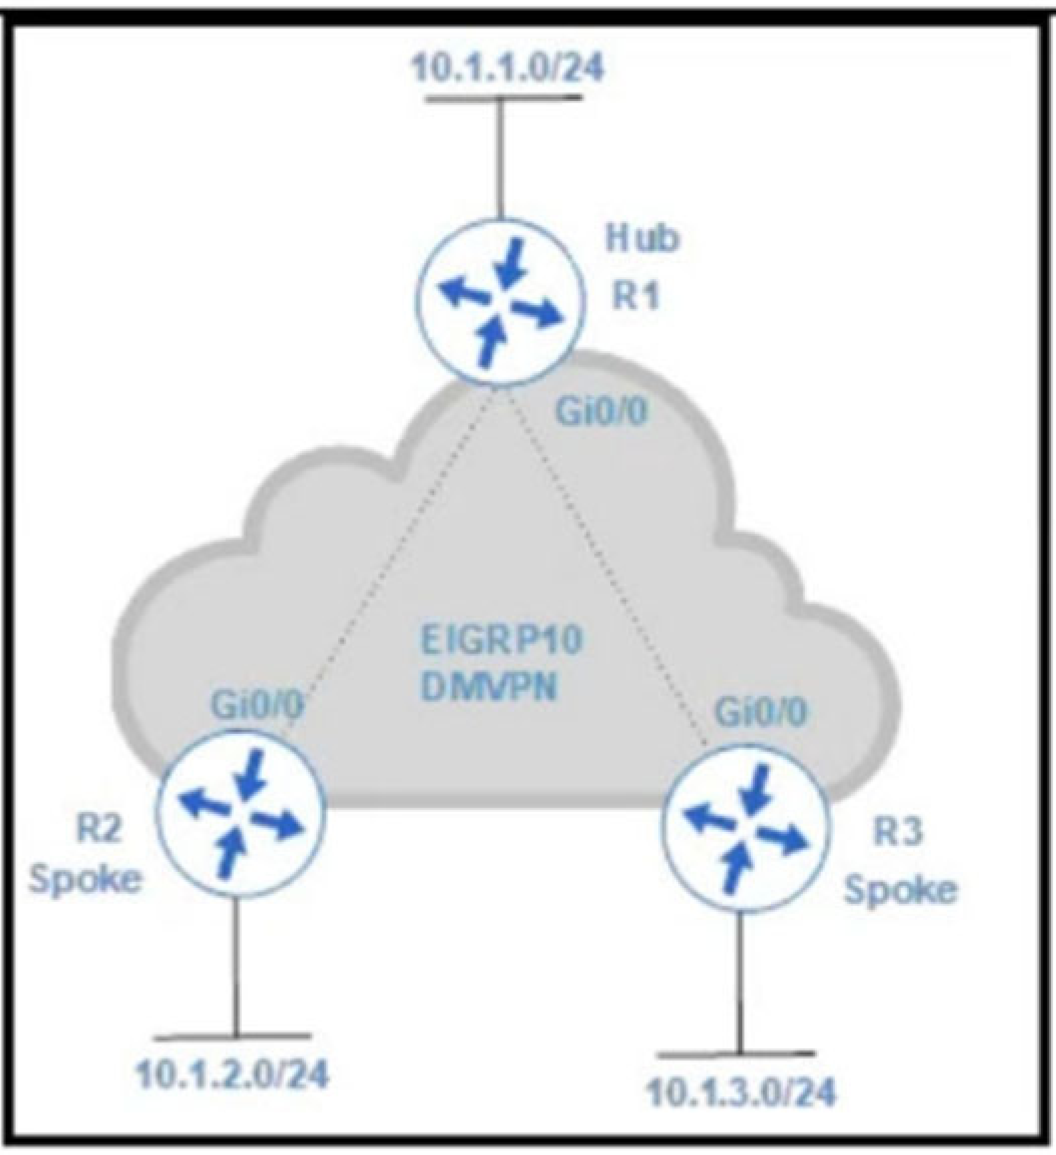

Refer to the exhibit. An engineer configures DMVPN and receives the hub location prefix of 10.1.1.0/24 on R2 and R3. The R3 prefix of 10.1.3.0/24 is not received on R2, and the R2 prefix 10.1.2.0/24 is not received on R3. Which action resolves the issue?

A. Split horizon prevents the routes from being advertised between spoke routers. It should be disabled with the no ip split-horizon eigrp 10 command on the Gi0/0 interface of R1.

B. There is no spoke-to-spoke connection. DMVPN configuration should be modified with a manual neighbor relationship configured between R2 and R3 and confirmed by use of the show ip eigrp neighbor command.

C. There is no spoke-to-spoke connection. DMVPN configuration should be modified to enable a tunnel connection between R2 and R3 and neighbor relationship confirmed by use of the show ip eigrp neighbor command.

D. Split horizon prevents the routes from being advertised between spoke routers. It should be disabled with the command no ip split-horizon eigrp 10 on the tunnel interface of R1.

Correct Answer: *D*

Community vote distribution

D (100%)

inteldarvid 5 months, 1 week ago

Selected Answer: D

D correct

upvoted 1 times

forccnp 10 months ago

Given answer in correct

.

upvoted 1 times

chris7890 1 year, 2 months ago

Answer D is correct: https://networkdirection.net/articles/routingandswitching/dmvpn/dmvpn-and-dynamic-routing/

upvoted 4 times

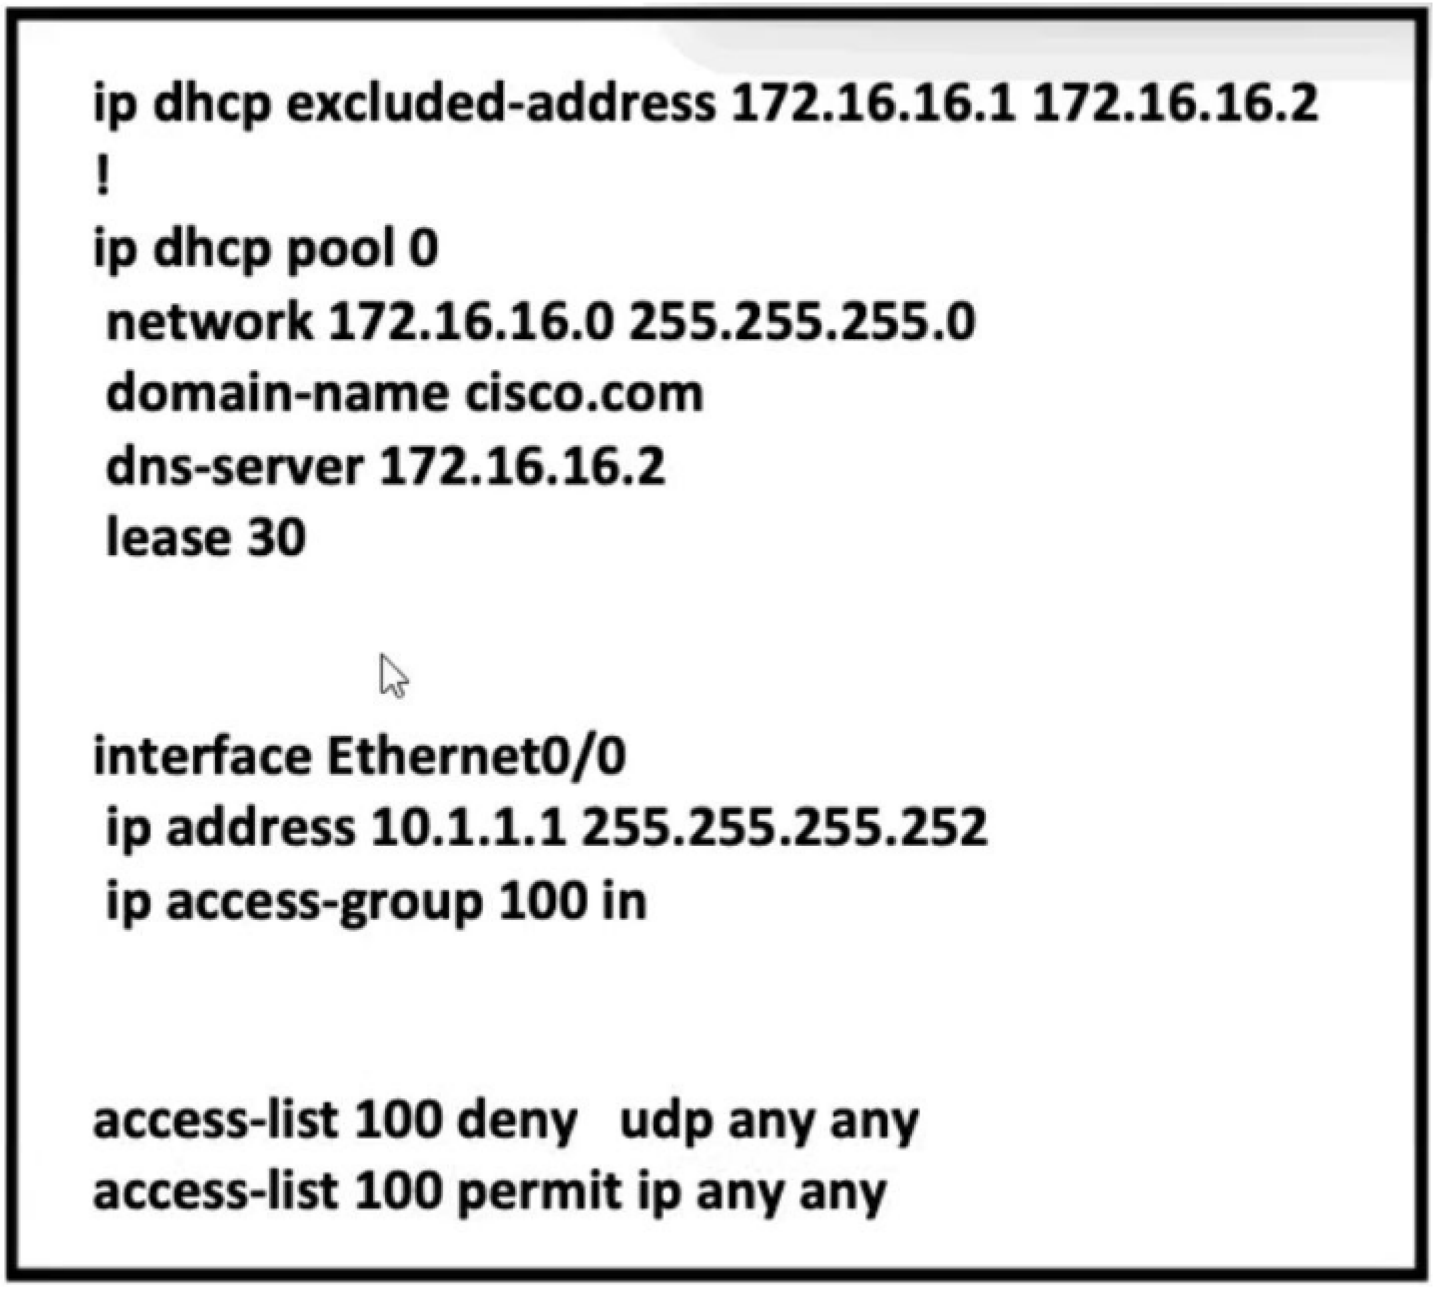

```
ip dhcp excluded-address 172.16.16.1 172.16.16.2
!
ip dhcp pool 0
 network 172.16.16.0 255.255.255.0
 domain-name cisco.com
 dns-server 172.16.16.2
 lease 30


interface Ethernet0/0
 ip address 10.1.1.1 255.255.255.252
 ip access-group 100 in


access-list 100 deny   udp any any
access-list 100 permit ip any any
```

Refer to the exhibit. Which two configurations allow clients to get dynamic IP addresses assigned? (Choose two.)

A. Configure access-list 100 permit udp any any eq 68 as the first line

B. Configure access-list 100 permit udp any any eq 69 as the first line

C. Configure access-list 100 permit udp any any eq 61 as the first line

D. Configure access-list 100 permit udp any any eq 66 as the first line

E. Configure access-list 100 permit udp any any eq 67 as the first line

Correct Answer: *AE*

GodFather Highly Voted 10 months, 3 weeks ago

DHCP servers have a UDP port number of 67
DHCP clients have the UDP port number 68

upvoted 8 times

guy276465281819372 Most Recent 5 months ago

Don't see why we need two answers, WE only need port 67 for server access.

upvoted 1 times

[Removed] 4 months, 1 week ago

you're right. The Client listens on 68 for the offer and ack parts of DORA, and because this ACL is inbound, port 68 should not matter as a destination port, it would matter as a source port.

upvoted 2 times

guy276465281819372 4 months ago

exactly we only need 67 here

upvoted 1 times

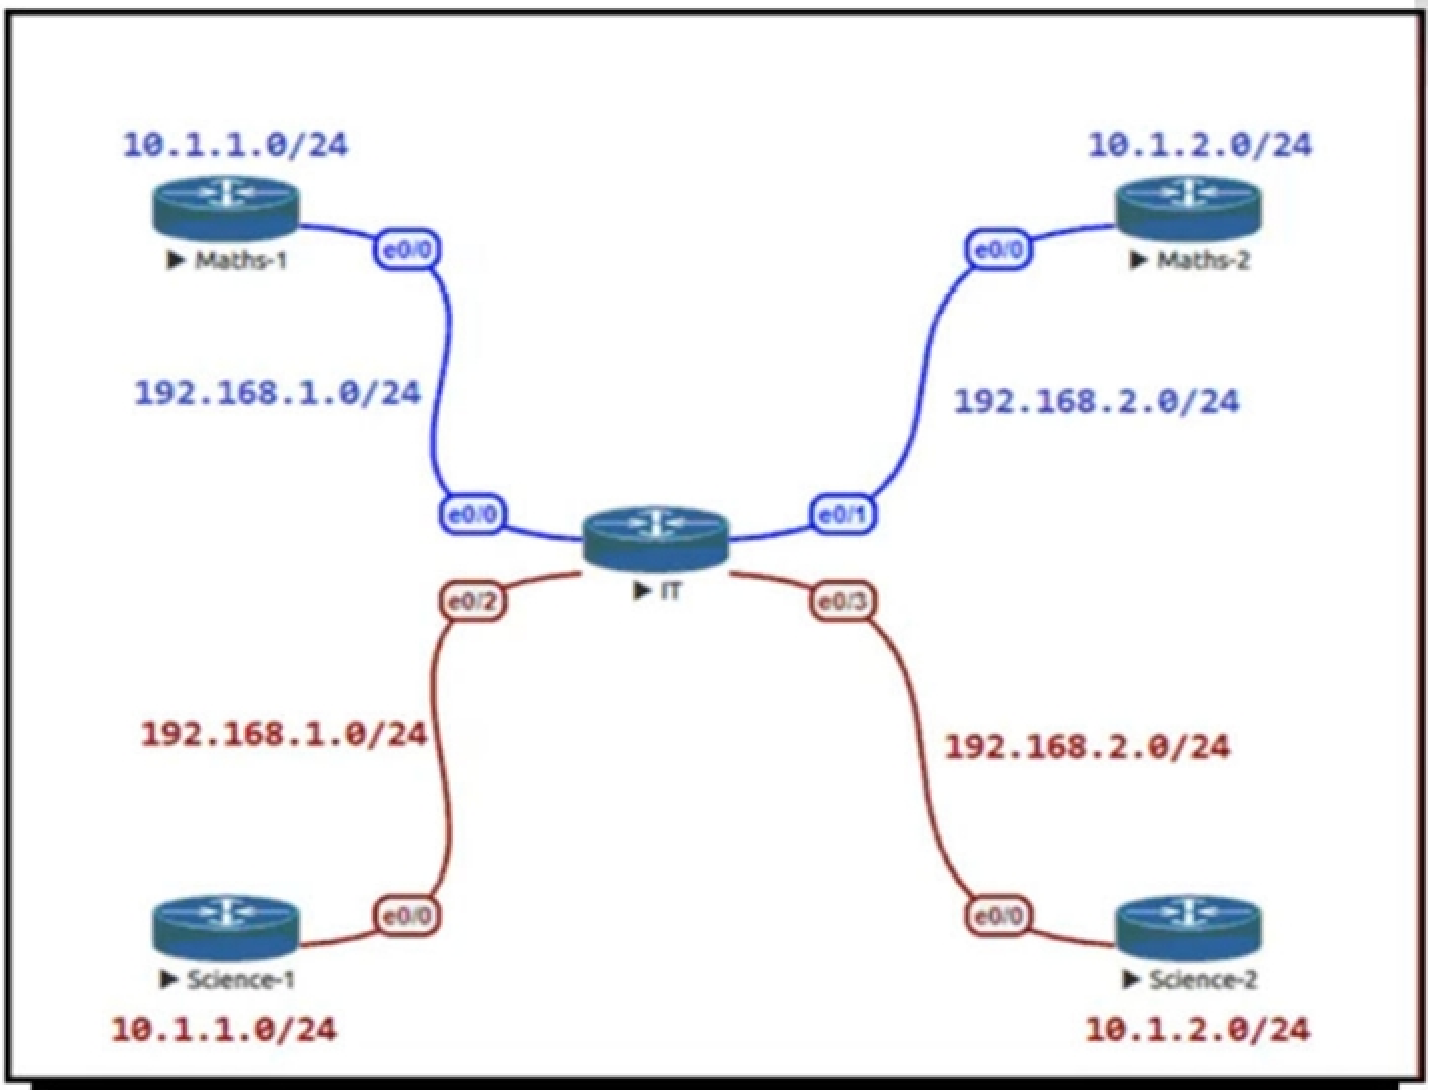

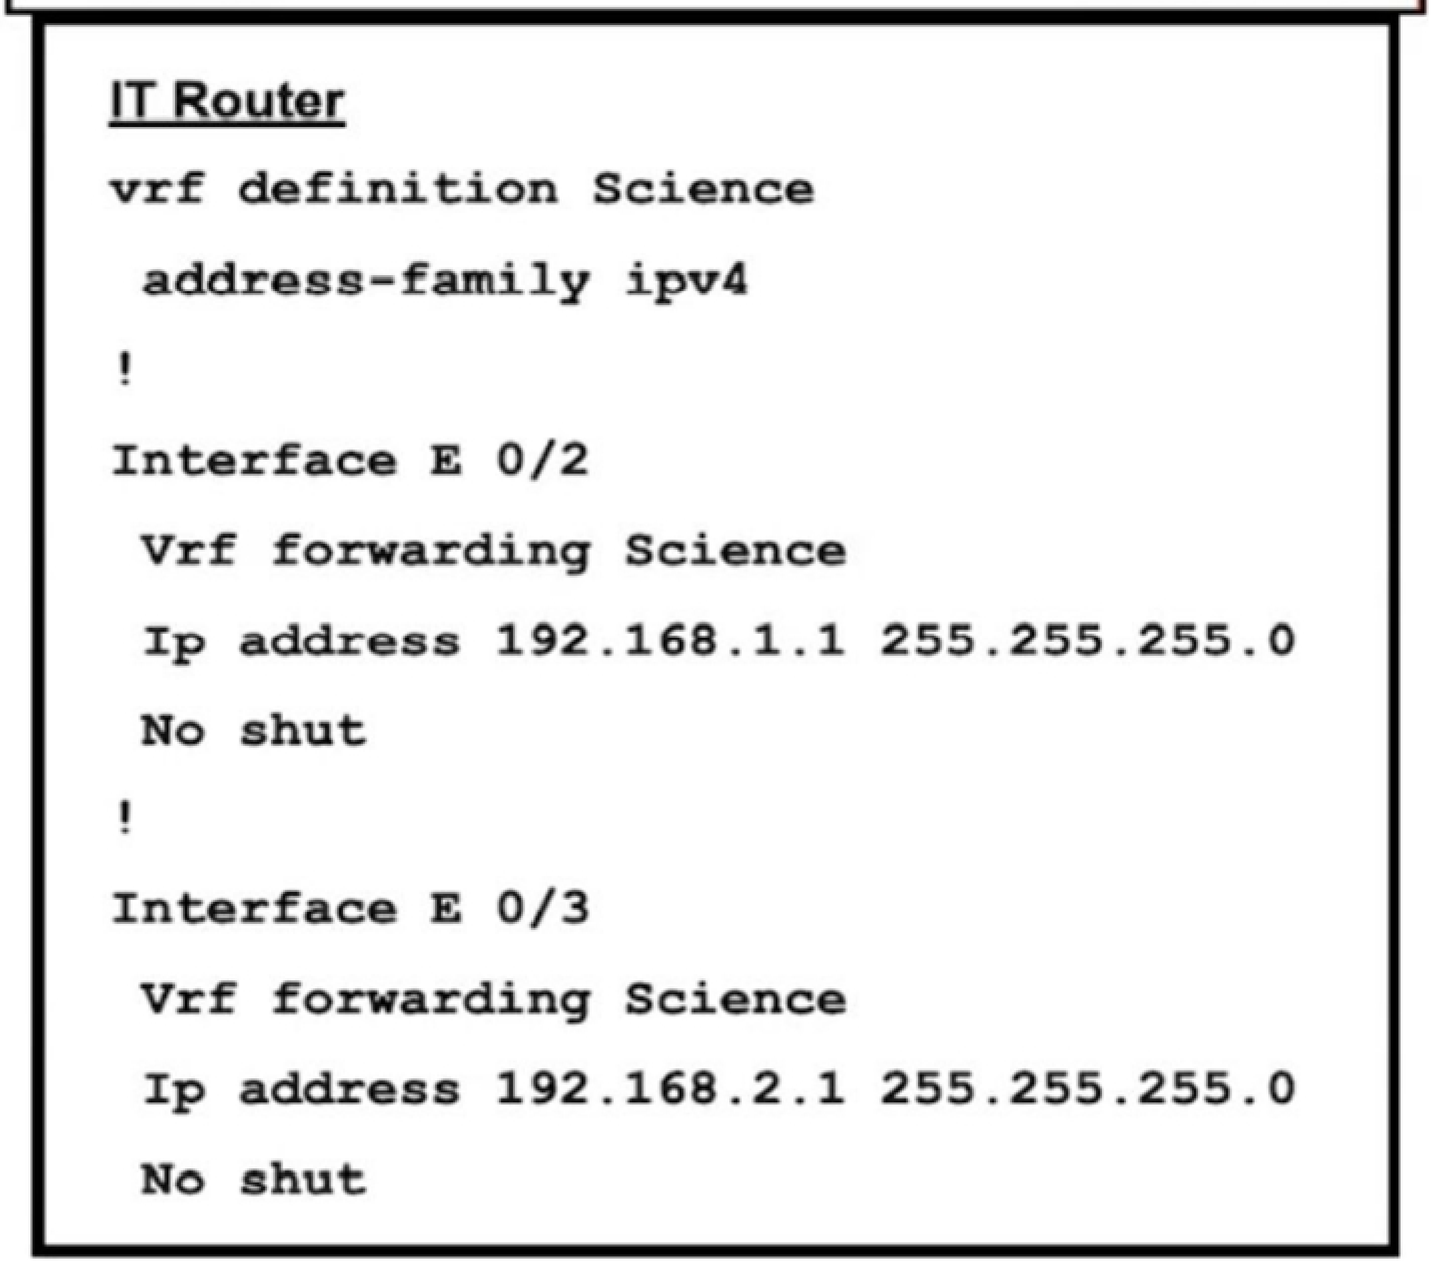

IT Router

```
vrf definition Science
 address-family ipv4
!
Interface E 0/2
 Vrf forwarding Science
 Ip address 192.168.1.1 255.255.255.0
 No shut
!
Interface E 0/3
 Vrf forwarding Science
 Ip address 192.168.2.1 255.255.255.0
 No shut
```

Refer to the exhibit. The IT router has been configured with the Science VRF and the interfaces have been assigned to the VRF. Which set of configurations advertises Science-1 and Science-2 routes using EIGRP AS 111?

A. router eigrp 111 address-family ipv4 vrf Science autonomous-system 1 network 192.168.1.0 network 192.168.2.0

B. router eigrp 111 address-family ipv4 vrf Science network 192.168.1.0 network 192.168.2.0

C. router eigrp 111 network 192.168.1.0 network 192.168.2.0

D. router eigrp 1 address-family ipv4 vrf Science autonomous-system 111 network 192.168.1.0 network 192.168.2.0

Correct Answer: *A*

Community vote distribution

D (92%) 8%

bolbolskanes Highly Voted 👍 12 months ago

D is the correct answer. please correct
in EIGRP named mode
R1(config)#router eigrp TEST (1 or 111 is just a name)
Cisco wanna trick us to make money

upvoted 10 times

Chiaretta Most Recent 🕒 5 months, 1 week ago

Selected Answer: D

D is the correct answer AS must be 111

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: D

sorry my anwser before was wrong. Th e option correct is "D". I test in my lab

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: A

100% answer correct "A"

upvoted 1 times

6dd4aa0 8 months, 2 weeks ago

Why can Answer A be right too?

upvoted 3 times

Cyril_the_Squirl 4 months, 1 week ago

A & D are perfectly correct....the question does require you to use AS 111, making D correct.

upvoted 1 times

Almylle 5 months, 3 weeks ago

Because the question asks for AS 111, not the process

upvoted 1 times

forccnp 9 months ago

Selected Answer: D

D is correct answer

upvoted 2 times

Typovy 9 months, 2 weeks ago

Just labbed it, if you will use answer B commands the AS for the vrf will be '0'. D is correct asnwer

upvoted 2 times

ChillingAgain 1 year, 1 month ago

Selected Answer: D

VRF-Lite for EIGRP using classic mode config.

upvoted 3 times

Huntkey 1 year, 2 months ago

Selected Answer: D

I like D too

upvoted 2 times

jarz 1 year, 2 months ago

Selected Answer: D

Ans = D

upvoted 2 times

lisanta12 1 year, 3 months ago

D is answer

upvoted 2 times

Question #276 *Topic 1*

An engineer must override the normal routing behavior of a router for Telnet traffic that is destined to 10.10.10.10 from 10.10.1.0/24 via a next hop of 10.4.4.4, which is directly connected to the router that is connected to the 10.1.1.0/24 subnet. Which configuration reroutes traffic according to this requirement?

A. access-list 100 deny tcp 10.10.1.0 0.0.0.255 host 10.10.10.10 eq 23 ! route-map POLICY permit 10 match ip address 100 set ip next-hop 10.4.4.4 route-map POLICY permit 20

B. access-list 100 permit tcp 10.10.1.0 0.0.0.255 host 10.10.10.10 eq 23 ! route-map POLICY permit 10 match ip address 100 set ip next-hop 10.4.4.4 route-map POLICY permit 20

C. access-list 100 permit tcp 10.10.1.0 0.0.0.255 host 10.10.10.10 eq 23 ! route-map POLICY permit 10 match ip address 100 set ip next-hop recursive 10.4.4.4 route-map POLICY permit 20

D. access-list 100 permit tcp 10.10.1.0 0.0.0.255 host 10.10.10.10 eq 23 ! route-map POLICY permit 10 match ip address 100 set ip next-hop recursive 10.4.4.4

Correct Answer: *C*

Community vote distribution

D (55%) B (40%) 5%

⊟ **VergilP** Highly Voted 👍 1 year, 1 month ago

Selected Answer: B

no need to config recursive

The recursive next-hop IP address is installed in the routing table and can be a subnet that is not directly connected. If the recursive next-hop IP address is not available, packets are routed using a default route.
---
https://www.cisco.com/en/US/docs/ios/iproute_pi/configuration/guide/iri_prb_rec_next_hop_external_docbase_0900e4b1810fe58b_4container_external_docbase_0900e4b181525fed.html

upvoted 12 times

⊟ **Patrick1234** Highly Voted 👍 10 months, 4 weeks ago

Selected Answer: D

I believe the 10.4.4.4 is not directly connected to this router, but is connected to a router behind 10.1.1.0/24 subnet. So recursive would be necessary. In that case I would go for answer D.

upvoted 9 times

⊟ **alex711** Most Recent ⏱ 3 months, 3 weeks ago

Selected Answer: D

D is correct.
route-map POLICY permit 20 is not used in PBR.
If you do not match packets on a route-map during PBR, PBR does not take any action on that packet, and is routed normally per the routing table/FIB/etc.

upvoted 1 times

⊟ **HarwinderSekhon** 4 months, 1 week ago

Selected Answer: B

There are 4 Devices 1. LAN PC 10.10.1.X/24 -- > Router directly connected to 10.10.1.X -->Router with IP 10.4.4.4 --> destination 10.10.10.10.

Just understand there are 4 nodes.
1.Client 10.10.1.X/24
2. Router connected to 10.10.1.X
3 Router we choose as next hop (10.4.4.4)
4. Destination 10.10.10.10
You are configuring node 2 and choosing node 3 as next hop. No recursive needed. Permit 20 does not harm in route map.

upvoted 2 times

⊟ **[Removed]** 4 months, 1 week ago

Selected Answer: D

D is the best answer.
At first I thought it was C, but I went back to my notes, a PBR does NOT require a second statement for traffic that is supposed to follow the RIB programming.
But Recursive keyword is required. Based on the wording of the problem it sounds like the router is not directly connected to 10.4.4.4.

"...override the normal routing behavior of a router...via next hop of 10.4.4.4 which is directly connected to the router that is connected to the 10.1.1.0/24 subnet..."

upvoted 2 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

team for me correct is "C", because the next hop (recursive) is remote and not connect directly and its necessary continue route map with seq "20", because block or deny rest traffic

upvoted 1 times

Almylle 5 months, 4 weeks ago

Selected Answer: D

For me D is the correct answer, because in this case u need the recursive command, the 10.4.4.4 is NOT directly connected to the router.

upvoted 2 times

Juraj22 5 months, 4 weeks ago

Selected Answer: C

draw a chema and you know that is not directly connected. Therefore must be recursive. co C or D, for me C is right, should be permit any at the end

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

I try to picture the path, but it's still not clear whether the "next-hop 10.4.4.4" is directly connected to the router with PBR or not.
Source: 10.10.1.0/24 || PBR || -> ??? -> next-hop 10.4.4.4 -> 10.1.1.0/24 -> destination: 10.10.10.10
B or D. Depends on the topology.

upvoted 4 times

6dd4aa0 8 months, 2 weeks ago

Selected Answer: B

B because it is directly connected, the option "recursive" does not need to be used.

upvoted 2 times

Typovy 9 months, 2 weeks ago

Selected Answer: D

Next hop router is connected to 10.1.1.0/24 but there is no info if it is directly connected to router on which we are configuring PBR. Since last permit is not needed in PBR the answer should be D

upvoted 3 times

Malasxd 7 months, 2 weeks ago

read the question again mate. they explicity say 10.4.4.4 is directly connected.

upvoted 1 times

JoeyT 6 months, 2 weeks ago

wrong. ... directly connected to the router that is connected to 10.1.1.0/24... means NOT directly connected...

upvoted 2 times

GReddy2323 9 months, 2 weeks ago

Why is permit 20 not needed in PBR?

upvoted 1 times

Typovy 9 months, 1 week ago

Ask cisco not me mate :D

upvoted 1 times

Titini 10 months, 1 week ago

Selected Answer: B

No need of recursive option

upvoted 1 times

Titini 10 months ago

Since it is not directly connected to this router C is the best option. We need also a permit 20 statement in the route map for the rest of the traffic. Sorry for the confusion.

upvoted 2 times

CisconAWSGURU 1 year, 1 month ago

Answer is B, Key work is "directly connected to the router"

upvoted 2 times

ChillingAgain 1 year, 1 month ago

Selected Answer: D

PBR does not need a default permit

upvoted 3 times

Huntkey 1 year, 2 months ago

Selected Answer: D

I never had to configure a default permit policy for PBR... I think D is good enough

upvoted 3 times

ChillingAgain 1 year, 1 month ago

Agree PBR does not need a default permit

upvoted 2 times

Question #277 Topic 1

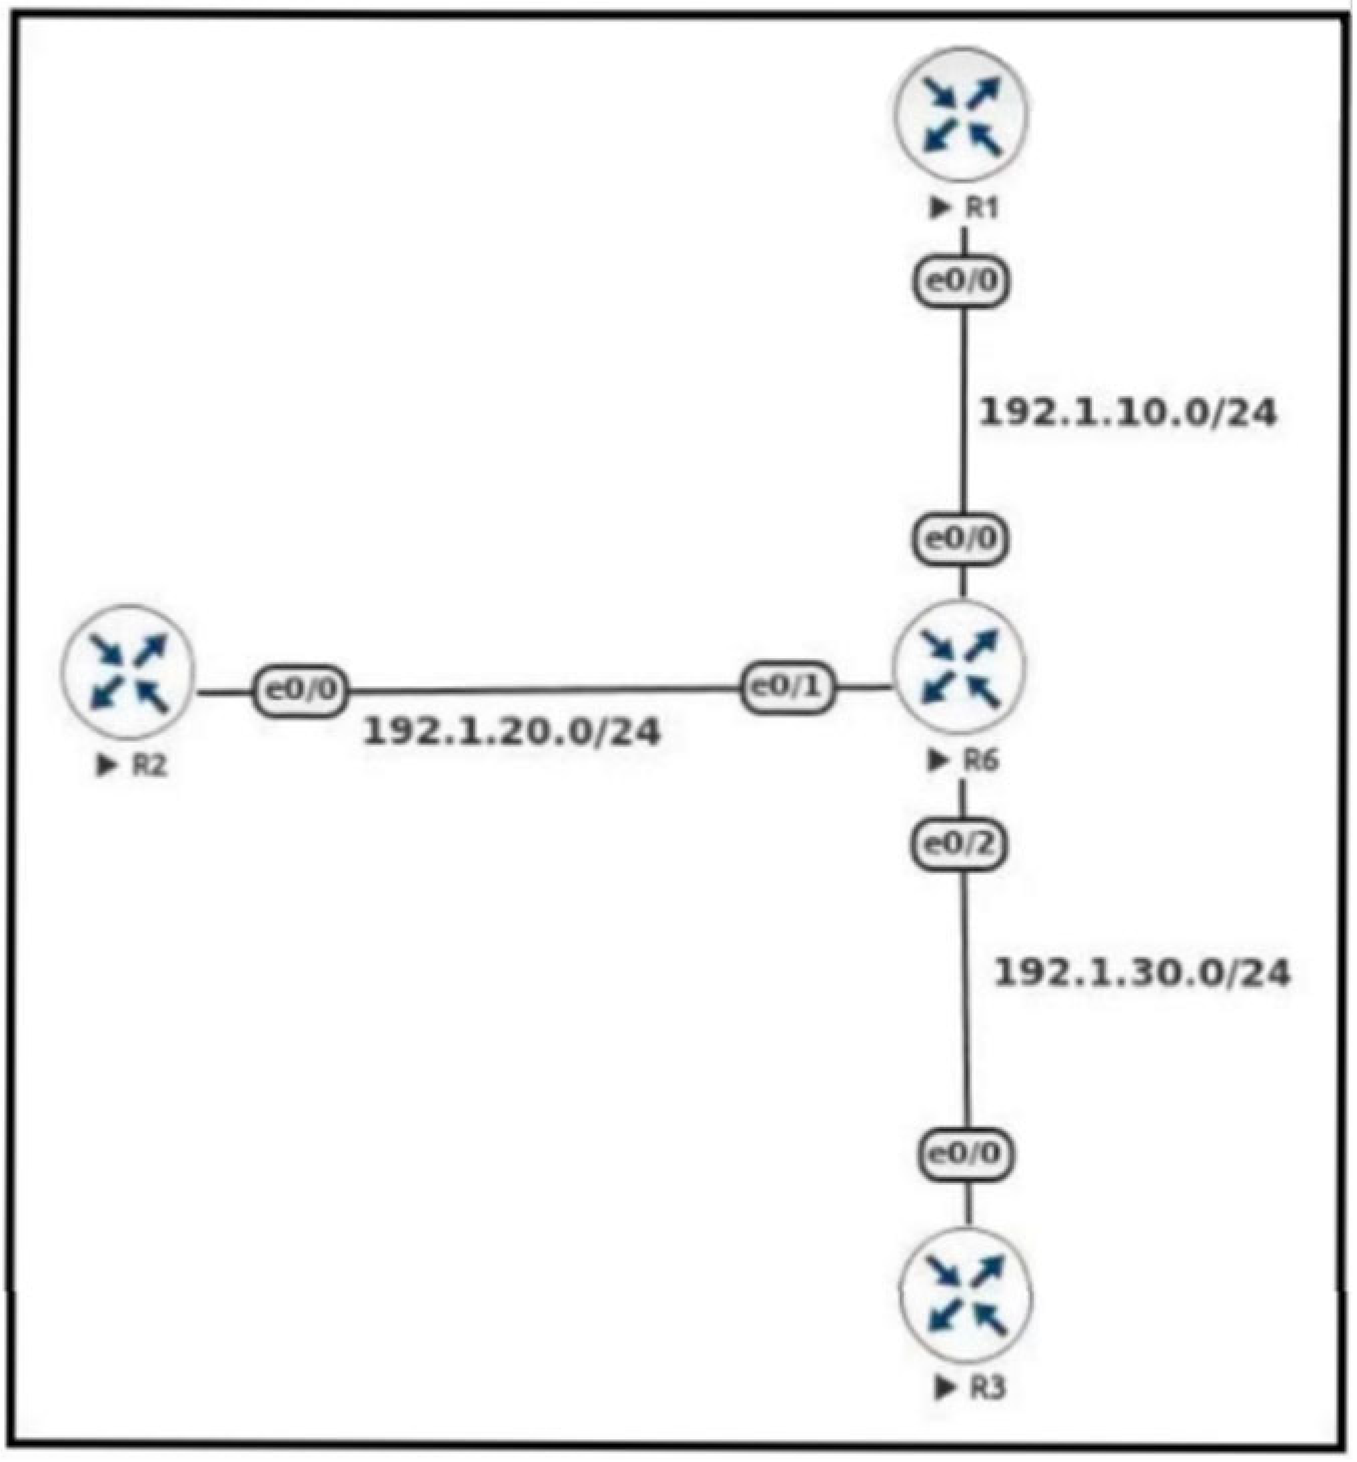

Refer to the exhibit. An engineer must configure DMVPN Phase 3 hub-and-spoke topology to enable a spoke-to-spoke tunnel. Which NHRP configuration meets the requirement on R6?

A. interface Tunnel1 ip nhrp authentication Cisco123 ip nhrp map multicast dynamic ip nhrp network-id 1 ip nhrp holdtime 300 ip nhrp redirect

B. interface Tunnel 1 ip address 192.168.1.1 255.255.255.0 tunnel source e 0/1 tunnel mode gre multipoint ip nhrp network-id 1 ip nhrp map 192.168.1.2 192.1.20.2

C. interface Tunnel1 ip nhrp authentication Cisco123 ip nhrp map multicast dynamic ip nhrp network-id 1 ip nhrp holdtime 300 ip nhrp shortcut

D. Interface Tunnel 1 ip address 192.168.1.1 255.255.255.0 tunnel source e 0/0 tunnel mode gre multipoint ip nhrp network-id 1

Correct Answer: *A*

Community vote distribution

A (100%)

Colmenarez 3 months, 3 weeks ago

Redirect is required on the hub

upvoted 1 times

HarwinderSekhon 4 months, 1 week ago

redirect on hub. shortcut command on spokes. DMVPN3

upvoted 2 times

[Removed] 4 months, 4 weeks ago

Is it not required to add the command “tunnel mode gre multipoint ” on the hub router

upvoted 1 times

HungarianDish 7 months, 2 weeks ago

Selected Answer: A

Only valid configuration on the Hub is "A".
https://ine.com/blog/2008-08-02-dmvpn-explained

upvoted 2 times

Question #278 Topic 1

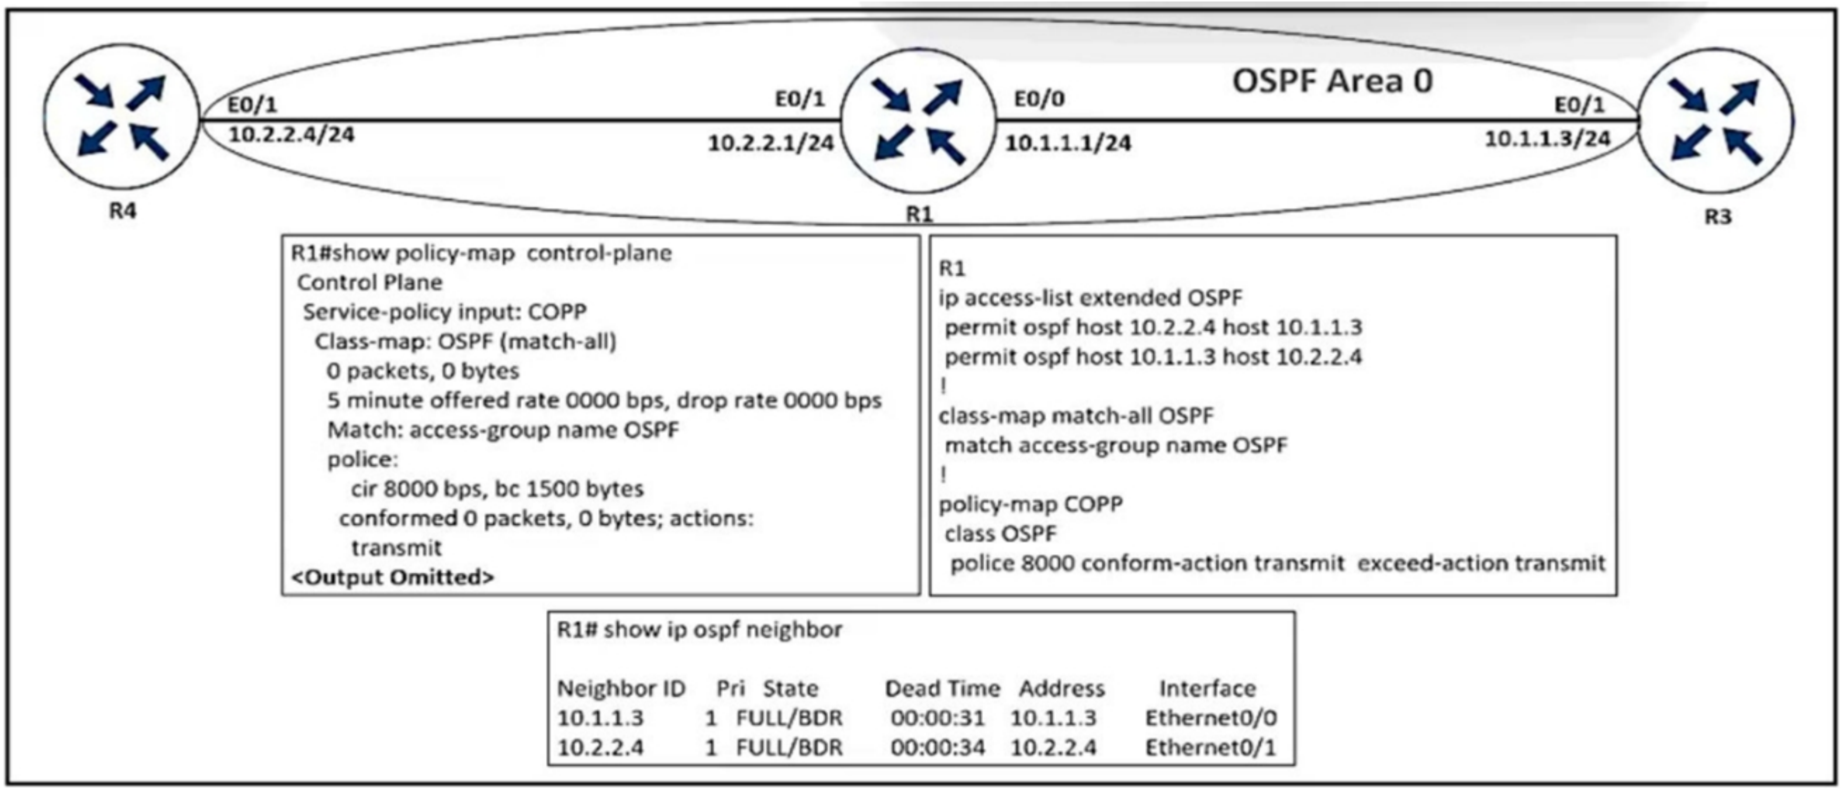

```
R1#show policy-map  control-plane
 Control Plane
  Service-policy input: COPP
    Class-map: OSPF (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: access-group name OSPF
      police:
          cir 8000 bps, bc 1500 bytes
        conformed 0 packets, 0 bytes; actions:
          transmit
<Output Omitted>
```

```
R1
ip access-list extended OSPF
 permit ospf host 10.2.2.4 host 10.1.1.3
 permit ospf host 10.1.1.3 host 10.2.2.4
!
class-map match-all OSPF
 match access-group name OSPF
!
policy-map COPP
 class OSPF
  police 8000 conform-action transmit  exceed-action transmit
```

```
R1# show ip ospf neighbor

Neighbor ID     Pri   State       Dead Time   Address      Interface
10.1.1.3          1   FULL/BDR    00:00:31    10.1.1.3     Ethernet0/0
10.2.2.4          1   FULL/BDR    00:00:34    10.2.2.4     Ethernet0/1
```

Refer to the exhibit. An engineer implemented CoPP but did not see OSPF traffic going through it. Which configuration resolves the issue?

A. control-plane service-policy input COPP

B. policy-map COPP class OSFP police 8000 conform-action transmit exceed-action transmit violate-action drop

C. ip access-list extended OSFP permit ospf any any

D. class-map match-all OSFP match access-group name OSFP

Correct Answer: *C*

Community vote distribution

C (77%) A (23%)

guy276465281819372 5 months ago

Selected Answer: C

definitely C

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

100 %%% is "C"

upvoted 1 times

HungarianDish 7 months, 2 weeks ago

Selected Answer: C

"A" and "D" are already applied, "B" is not required as traffic only needs to be captured and not limited, so "drop" is incorrect. "C" is correct, however it would be enough to set the appropriate source and destination IP pairs, as Aikat and others wrote.

upvoted 3 times

HungarianDish 7 months, 2 weeks ago

https://community.cisco.com/t5/switching/ospf-dies-when-apply-acl/td-p/794381

This thread suggested to use "permit ospf any any" for simplicity, because of the multicast addressing.

upvoted 1 times

Mikedask 9 months, 2 weeks ago

if the acl wasnt right then why we have full ospf adj?....i mean we have hellos exhange full/bdr and right ospf process.
if we hasnt full state then the right answer will be the c but i think the copp policy must be configured A

upvoted 1 times

yefrimart 2 months, 2 weeks ago

Remember than when the traffic do not match the policy it simply does not apply the policy and the traffic is treated normally. That is why we have full adjacencies between the routers.

upvoted 1 times

Aikat 9 months, 3 weeks ago

pay attention to IP pairs:
- 10.2.2.4 <> 10.2.2.1
- 10.1.1.1 <> 10.1.1.3

then check what's allowed in the ACL. Answer is C

upvoted 1 times

MD_Shox 1 year ago

Selected Answer: C

this is mcast and in addition look carefully at R1 R2 R3 interface ip addresses
only C can solve it from the listed answers and will catch bot R1<-R2 and R2<->R3

upvoted 2 times

VergilP 1 year, 1 month ago

Selected Answer: C

seems like ..... should be C because of the multicast..
I'm not very sure but I vote for C

https://community.cisco.com/t5/routing/access-list-ospf/td-p/781095

upvoted 1 times

Edwinmolinab 1 year, 1 month ago

Selected Answer: C

Given answer is correct tested on GNS3

upvoted 2 times

chris7890 1 year, 3 months ago

Selected Answer: A

The configured policy map must be assigned in the control plan
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/15-0SY/configuration/guide/15_0_sy_swcg/control_plane_policing_copp.pdf

upvoted 3 times

VergilP 1 year, 1 month ago

why? i see the COPP is already config?
the first two line of the left picture...

upvoted 2 times

Question #279

Topic 1

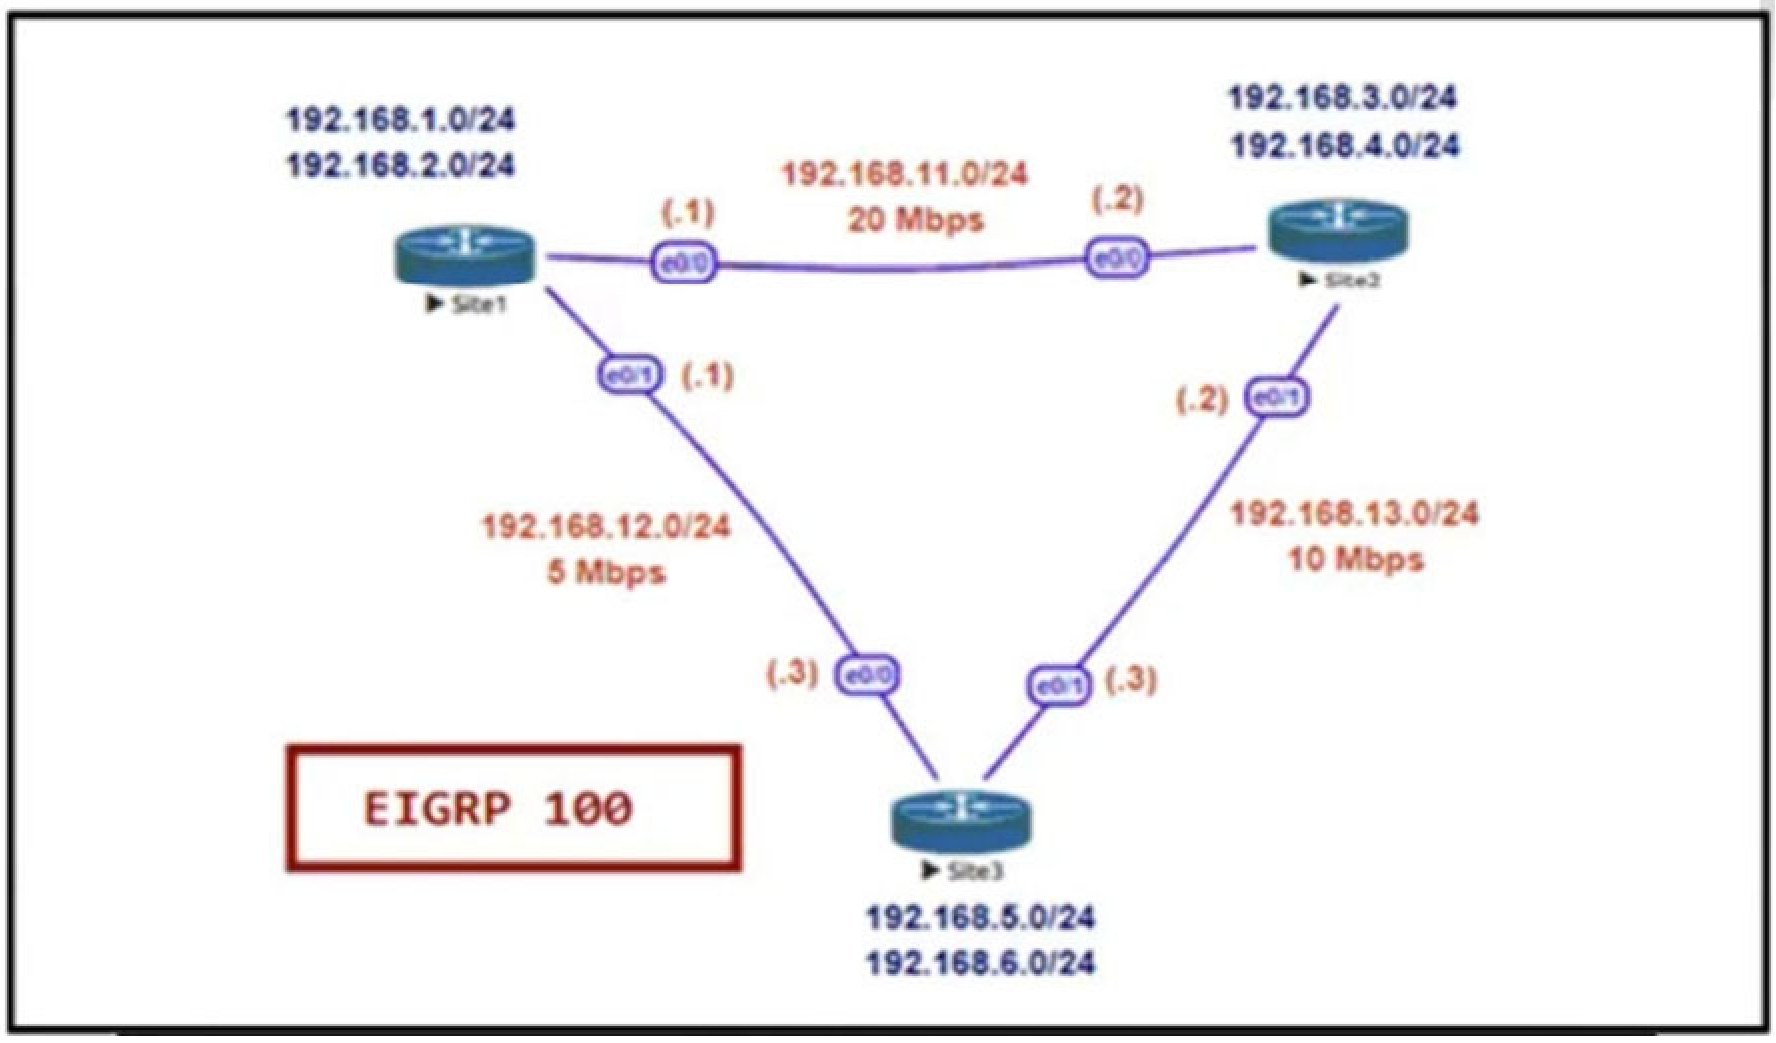

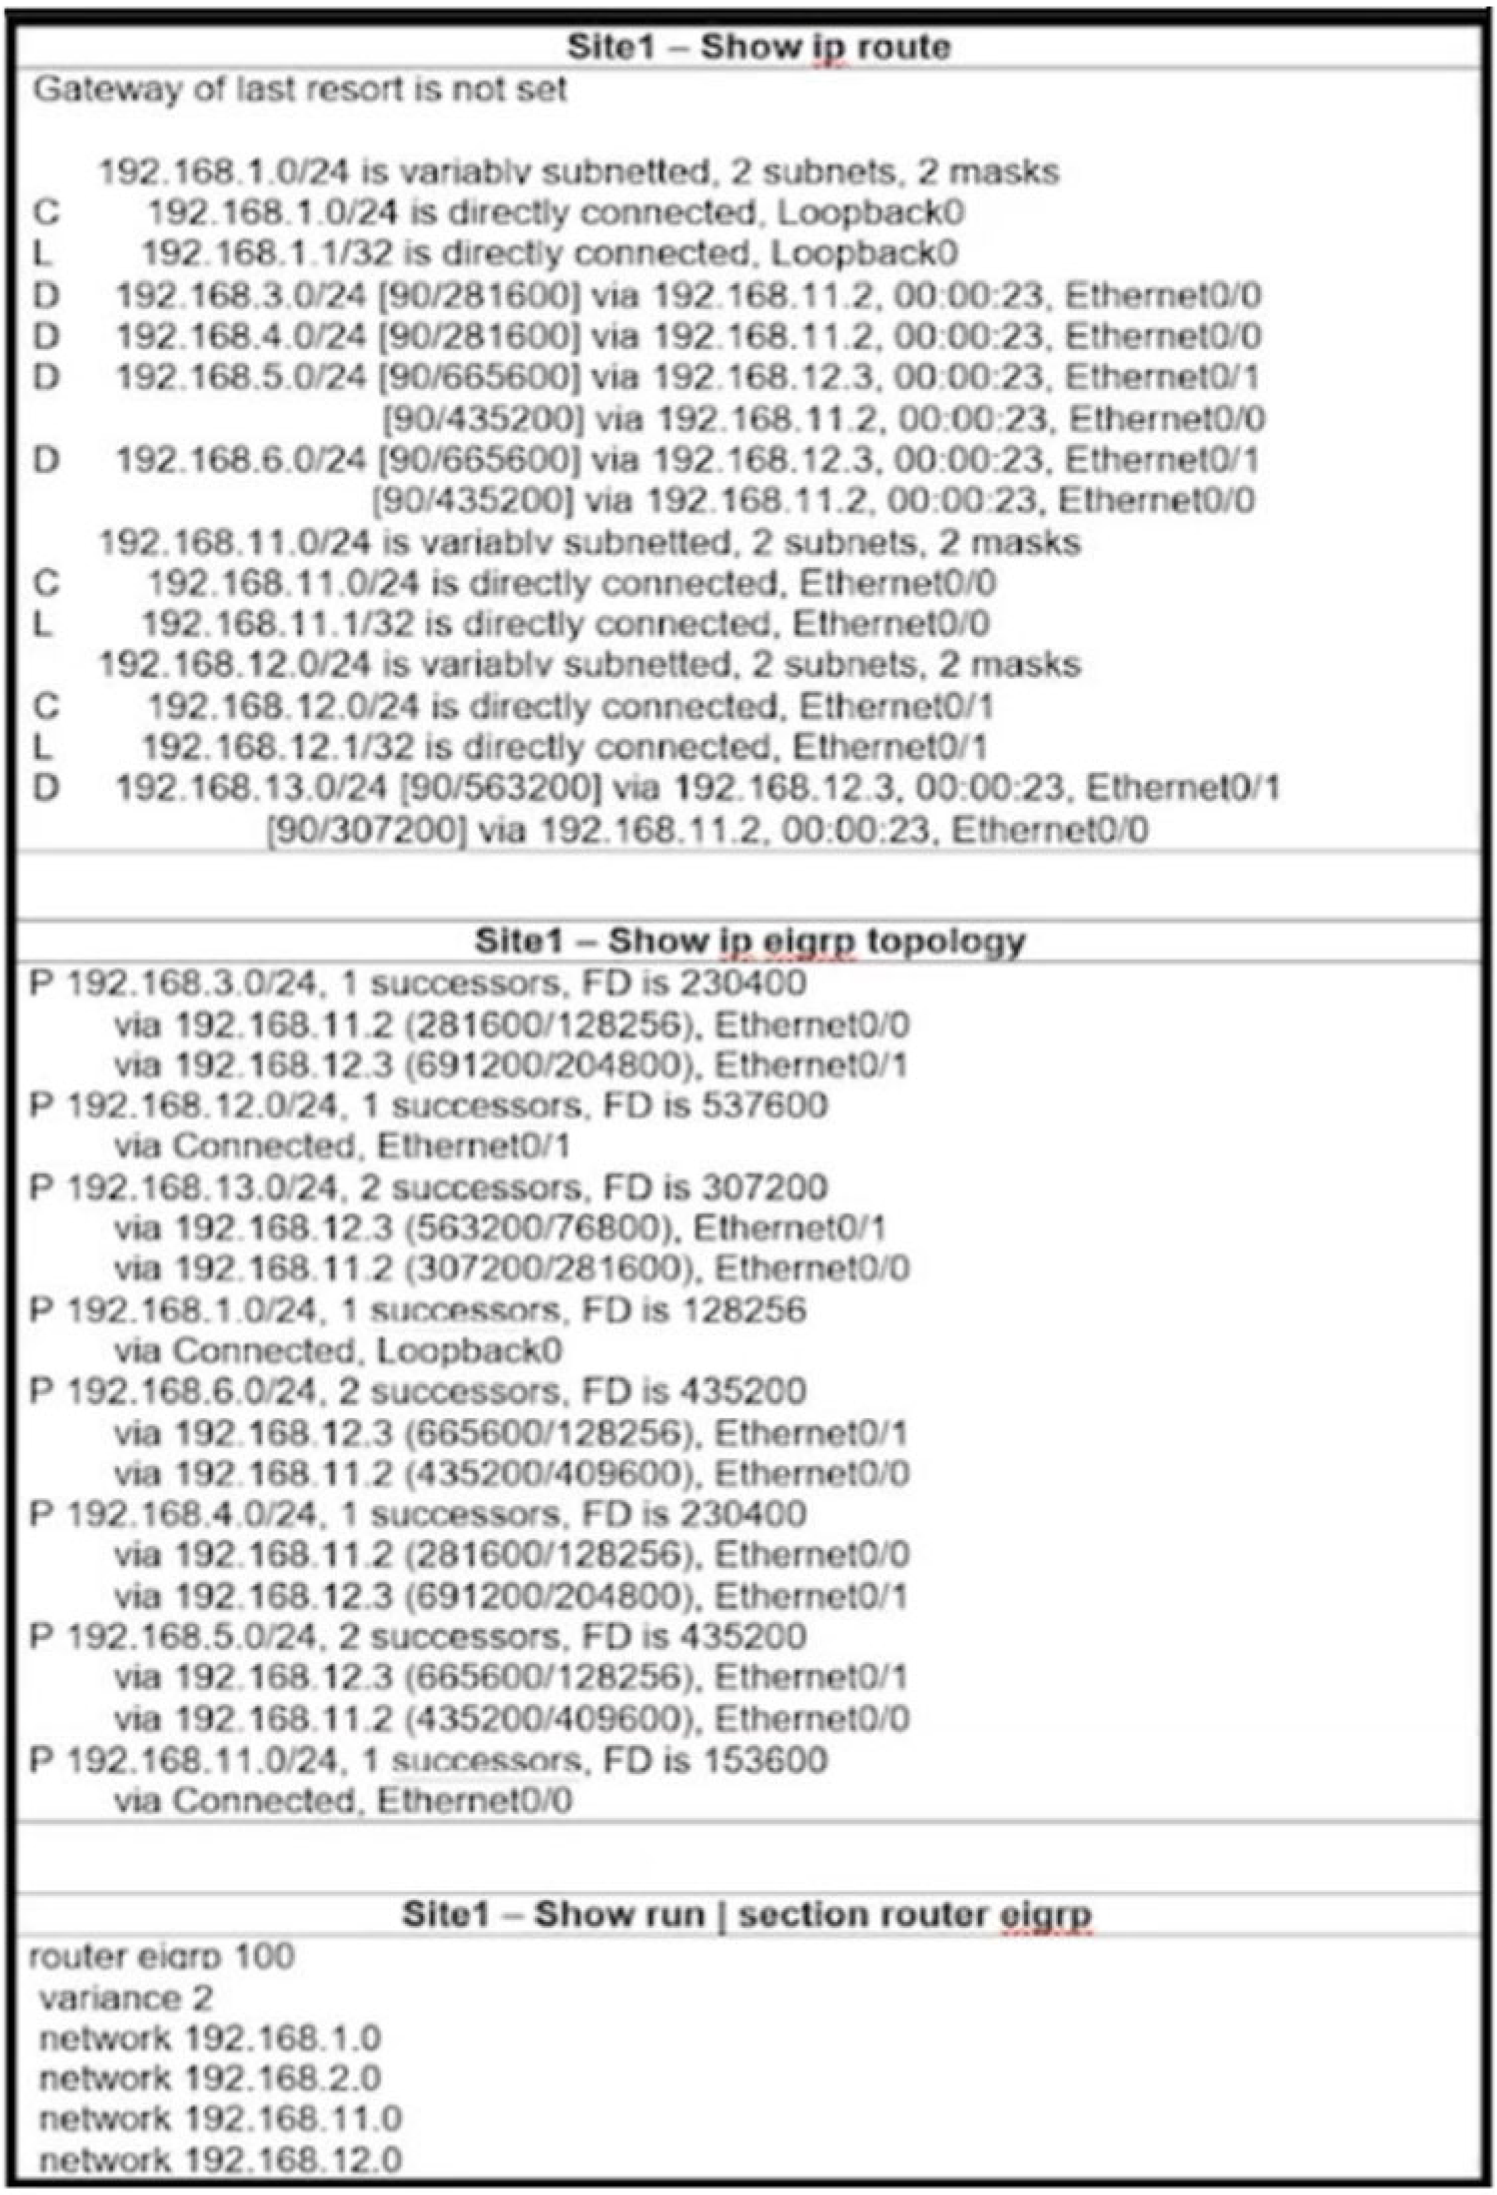

Site1 – Show ip route

```
Gateway of last resort is not set

     192.168.1.0/24 is variablv subnetted, 2 subnets, 2 masks
C       192.168.1.0/24 is directly connected, Loopback0
L       192.168.1.1/32 is directly connected, Loopback0
D    192.168.3.0/24 [90/281600] via 192.168.11.2, 00:00:23, Ethernet0/0
D    192.168.4.0/24 [90/281600] via 192.168.11.2, 00:00:23, Ethernet0/0
D    192.168.5.0/24 [90/665600] via 192.168.12.3, 00:00:23, Ethernet0/1
                    [90/435200] via 192.168.11.2, 00:00:23, Ethernet0/0
D    192.168.6.0/24 [90/665600] via 192.168.12.3, 00:00:23, Ethernet0/1
                    [90/435200] via 192.168.11.2, 00:00:23, Ethernet0/0
     192.168.11.0/24 is variablv subnetted, 2 subnets, 2 masks
C       192.168.11.0/24 is directly connected, Ethernet0/0
L       192.168.11.1/32 is directly connected, Ethernet0/0
     192.168.12.0/24 is variablv subnetted, 2 subnets, 2 masks
C       192.168.12.0/24 is directly connected, Ethernet0/1
L       192.168.12.1/32 is directly connected, Ethernet0/1
D    192.168.13.0/24 [90/563200] via 192.168.12.3, 00:00:23, Ethernet0/1
                [90/307200] via 192.168.11.2, 00:00:23, Ethernet0/0
```

Site1 – Show ip eigrp topology

```
P 192.168.3.0/24, 1 successors, FD is 230400
     via 192.168.11.2 (281600/128256), Ethernet0/0
     via 192.168.12.3 (691200/204800), Ethernet0/1
P 192.168.12.0/24, 1 successors, FD is 537600
     via Connected, Ethernet0/1
P 192.168.13.0/24, 2 successors, FD is 307200
     via 192.168.12.3 (563200/76800), Ethernet0/1
     via 192.168.11.2 (307200/281600), Ethernet0/0
P 192.168.1.0/24, 1 successors, FD is 128256
     via Connected, Loopback0
P 192.168.6.0/24, 2 successors, FD is 435200
     via 192.168.12.3 (665600/128256), Ethernet0/1
     via 192.168.11.2 (435200/409600), Ethernet0/0
P 192.168.4.0/24, 1 successors, FD is 230400
     via 192.168.11.2 (281600/128256), Ethernet0/0
     via 192.168.12.3 (691200/204800), Ethernet0/1
P 192.168.5.0/24, 2 successors, FD is 435200
     via 192.168.12.3 (665600/128256), Ethernet0/1
     via 192.168.11.2 (435200/409600), Ethernet0/0
P 192.168.11.0/24, 1 successors, FD is 153600
     via Connected, Ethernet0/0
```

Site1 – Show run | section router eigrp

```
router eiarp 100
 variance 2
 network 192.168.1.0
 network 192.168.2.0
 network 192.168.11.0
 network 192.168.12.0
```

Refer to the exhibit. Site1 must perform unequal cost load balancing toward the segments behind Site2 and Site3. Some of the routes are getting load balanced but others are not. Which configuration allows Site1 to load balance toward all the LAN segments of the remote routers?

A. Site3 router eigrp 100 variance 2

B. Site2 router eigrp 100 variance 2

C. Site2 router eigrp 100 variance 3

D. Site1 router eigrp 100 variance 3

Correct Answer: *D*

Community vote distribution

D (100%)

Question #280

Topic 1

ıt, then

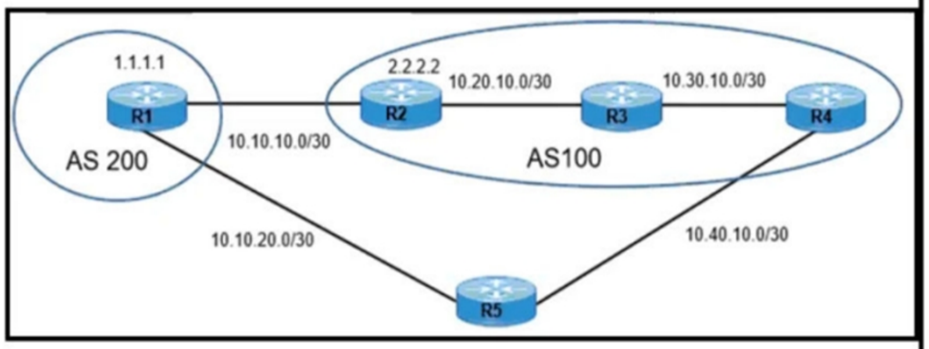

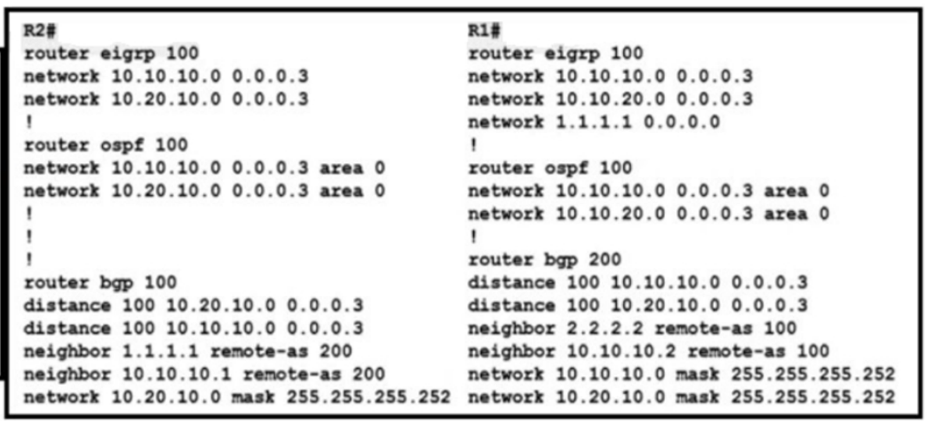

```
R2#                                         R1#
router eigrp 100                            router eigrp 100
network 10.10.10.0 0.0.0.3                  network 10.10.10.0 0.0.0.3
network 10.20.10.0 0.0.0.3                  network 10.10.20.0 0.0.0.3
!                                           network 1.1.1.1 0.0.0.0
router ospf 100                             !
network 10.10.10.0 0.0.0.3 area 0           router ospf 100
network 10.20.10.0 0.0.0.3 area 0           network 10.10.10.0 0.0.0.3 area 0
!                                           network 10.10.20.0 0.0.0.3 area 0
!                                           !
!                                           router bgp 200
router bgp 100                              distance 100 10.10.10.0 0.0.0.3
distance 100 10.20.10.0 0.0.0.3             distance 100 10.20.10.0 0.0.0.3
distance 100 10.10.10.0 0.0.0.3             neighbor 2.2.2.2 remote-as 100
neighbor 1.1.1.1 remote-as 200              neighbor 10.10.10.2 remote-as 100
neighbor 10.10.10.1 remote-as 200           network 10.10.10.0 mask 255.255.255.252
network 10.20.10.0 mask 255.255.255.252     network 10.20.10.0 mask 255.255.255.252
```

Refer to the exhibit. R1 and R2 use IGP protocol to route traffic between AS 100 and AS 200 despite being configured to use BGP. Which action resolves the issue and ensures the use of BGP?

A. Configure distance to 100 under the OSPF process of R1 and R2

B. Remove distance commands under BGP AS 100

C. Remove distance commands under BGP AS 100 and AS 200.

D. Configure distance to 100 under the EIGRP process of R1 and R2

Correct Answer: *B*

Community vote distribution

C (100%)

inteldarvid 5 months, 1 week ago

Selected Answer: C

100 %% option "C"

upvoted 1 times

Almylle 5 months, 4 weeks ago

Selected Answer: C

I think the aswer is C, but only removing the distance 100 between neighbors 100 and 200

upvoted 1 times

forccnp 9 months, 1 week ago

Selected Answer: C

C is correct answer,
Remove distance 100 command from both router

upvoted 2 times

bolbolskanes 12 months ago

The question is unclear
eBGP = 20 so it will be the preferred

upvoted 1 times

ellen_AA 11 months, 2 weeks ago

But its AD is overwritten to 100 using distance command. Removing the distance command brings eBGP Ad back to 20. So BGP will be installed in the routing table.

upvoted 5 times

MD_Shox 1 year ago

Selected Answer C

upvoted 1 times

Noproblem22 1 year ago

Why only under AS 100? I think the correct answer is C

upvoted 1 times

SDWAN 1 year, 2 months ago

C. Take BGP 100 out... Ebgp 20 is preferred.

upvoted 1 times

Huntkey 1 year, 2 months ago

I don't know any answer is correct... Even after changing the AD for BGP, it is still better than the OSPF AD of 110...

upvoted 3 times

chris7890 1 year, 3 months ago

Selected Answer: C

I think the correct answer is C. Remove removal commands under BGP AS 100 and AS 200.

upvoted 4 times

shoo83 11 months, 3 weeks ago

R1 & R2 establish EIGRP on segment 10.10.10.0/30

upvoted 2 times

Question #281 *Topic 1*

DRAG DROP -

Drag and drop the MPLS concepts from the left onto the descriptions on the right.

Select and Place:

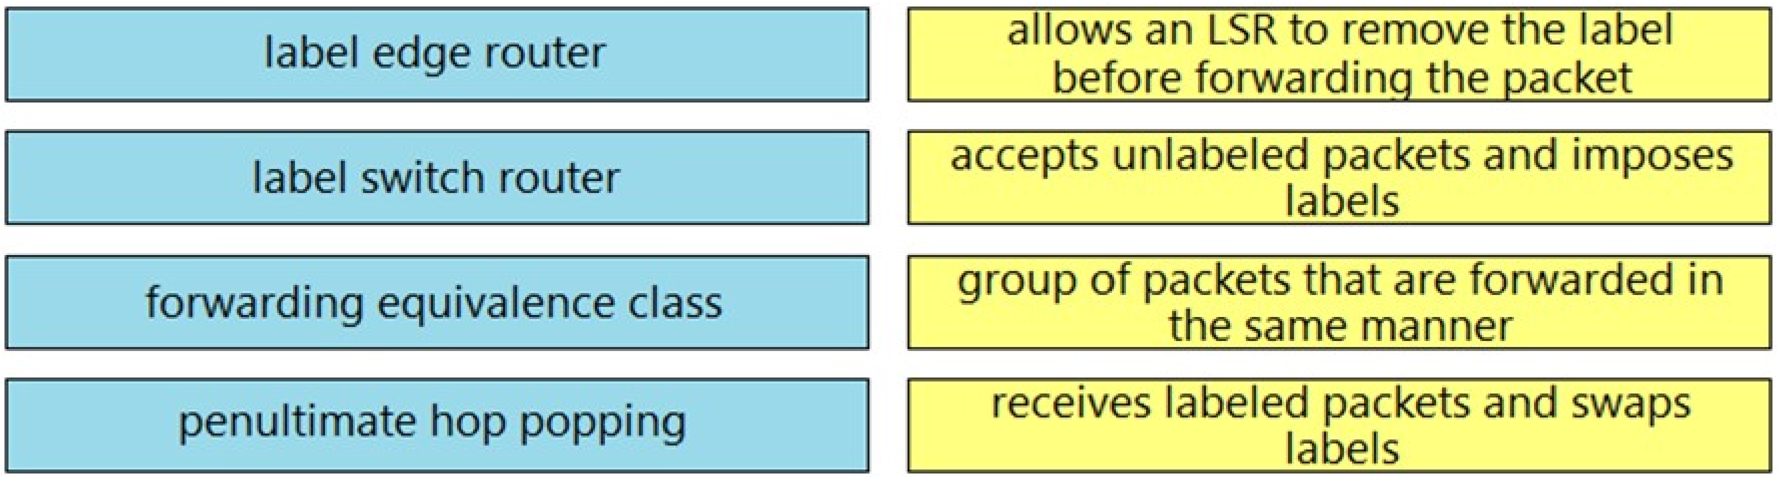

label edge router	allows an LSR to remove the label before forwarding the packet
label switch router	accepts unlabeled packets and imposes labels
forwarding equivalence class	group of packets that are forwarded in the same manner
penultimate hop popping	receives labeled packets and swaps labels

Correct Answer:

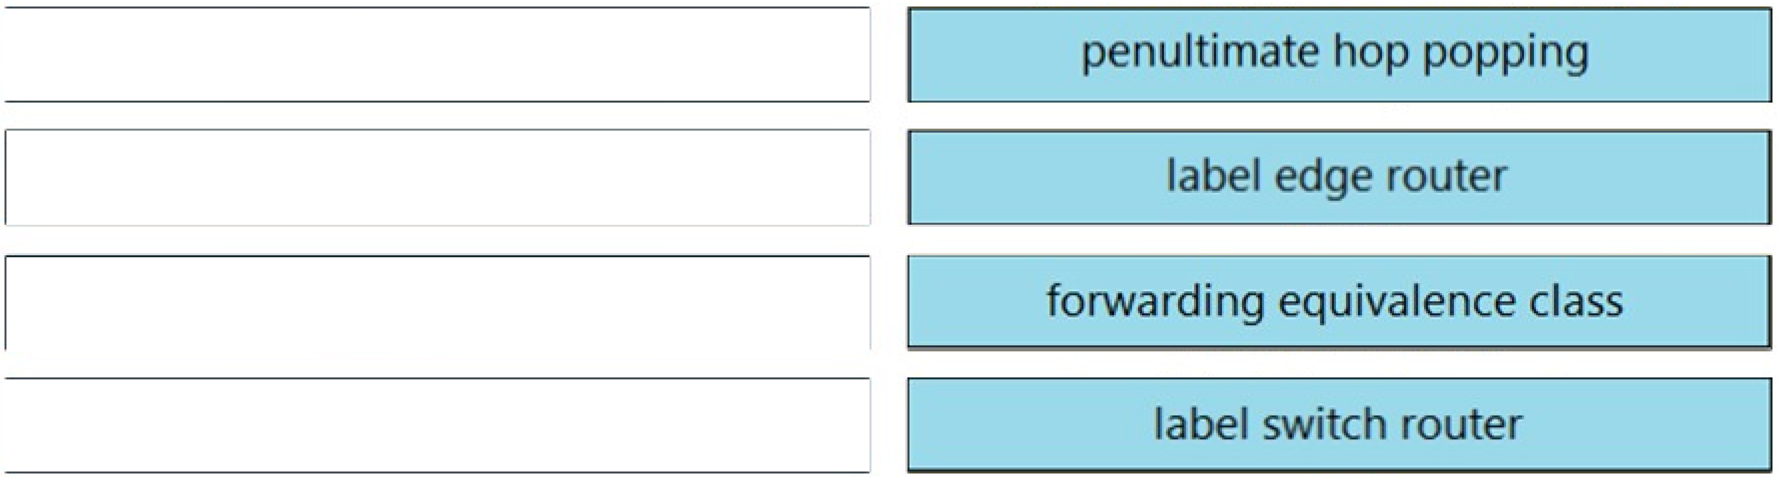

	penultimate hop popping
	label edge router
	forwarding equivalence class
	label switch router

inteldarvid 5 months, 1 week ago

correct

upvoted 1 times

Xerath 9 months, 4 weeks ago

The given answer is correct.

upvoted 2 times

Question #282 *Topic 1*

Which table is used to map the packets in an MPLS LSP that exit from the same interface, via the same next hop, and have the same queuing policies?

A. LDP

B. FEC

C. CEF

D. RIB

Correct Answer: *B*

Community vote distribution

B (67%) C (33%)

ZamanR 6 days, 6 hours ago

B is the Answer

upvoted 1 times

Brand 3 months, 2 weeks ago

According to the link HungarianDish provided, it seems they are asking for CEF as it is an actual "table" use to populate FEC attributes.

upvoted 1 times

JieW 4 months, 2 weeks ago

Selected Answer: B

CEF isnt a table either. My guess is FEC.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: B

FOR ME IS "B", Because, I think there is a problem in the question, the word "table", the rest of the question is the same FEC concept, the same label for several pefixes with the same next hop and the same queuing policies.

upvoted 2 times

HungarianDish 7 months, 2 weeks ago

Selected Answer: C

For me it's CEF, because this table is used for creating the LSP. Plus, FEC is not a table, it is rather an attribute (e.g. a destination IP subnet is a typical FEC).

upvoted 2 times

HungarianDish 7 months, 2 weeks ago

https://community.cisco.com/t5/routing/why-cef-needed-in-mpls-network/td-p/1699091
cisco MPLS code ...uses as input data the FIB (Forwarding Information Base) mantained by CEF,
to build the LFIB that is the table where for each FEC there is an association with a label taken from the local node label space.
...the biggest difference is that the CEF table is kept local and not exported to any other device. MPLS FEC/label bindings are advertised.

upvoted 1 times

HungarianDish 7 months, 2 weeks ago

https://www.networkworld.com/article/2291724/chapter-7--understanding-cef-in-an-mpls-vpn-environment.html
MPLS creates its own database for lookups called the Label Forwarding Information Base (LFIB),
but it uses the CEF FIB as a source of this information.
In the direction of label imposition, the router switches packets based on a CEF table lookup to find the next hop
and adds the appropriate label information stored in the FIB for the destination.

upvoted 1 times

HungarianDish 7 months, 2 weeks ago

The question describes FEC, however, the table which being used is Label Forwarding Information Base (LFIB) in Cisco terms or "FEC-to-NHLFE" (FTN) table according to RFC 3031. LFIB is using CEF table + LIB.

As well as I see, none of the answers are correct.

upvoted 1 times

HungarianDish 7 months, 2 weeks ago

Good explanations:
https://community.cisco.com/t5/routing/mpls-tables/td-p/2305490

https://www.ccexpert.us/routing-switching/mpls-packet-forwarding-and-label-switched-paths.html

upvoted 1 times

HungarianDish 7 months, 2 weeks ago

RFC 3031
https://datatracker.ietf.org/doc/rfc3031/

upvoted 1 times

msama 1 year, 1 month ago

Selected Answer: B

A forwarding equivalence class (FEC) is a term used in Multiprotocol Label Switching (MPLS) to describe a set of packets with similar or identical characteristics which may be forwarded the same way; that is, they may be bound to the same MPLS label.

upvoted 1 times

IceFireSoul 1 year, 2 months ago

Given Answer is correct, for references see:
https://learningnetwork.cisco.com/s/question/0D53i00000Ksx8ZCAR/what-is-fec-in-mpls-and-how-it-works-

upvoted 2 times

Question #283 *Topic 1*

You have configured router R1 with multiple VRF's in order to support multiple customer VPN networks. If you wanted to see the best path for the 10.1.1.0.24 route in VRF Blue, what command would you use?

A. show ip route vrf Blue 10.1.1.0

B. show ip route 10.1.1.0 vrf Blue

C. show route all 10.1.1.0

D. show ip route all 10.1.1.0

Correct Answer: *A*

Community vote distribution

A (100%)

HungarianDish 7 months, 2 weeks ago

Selected Answer: A

Agree with the answer

upvoted 2 times

Question #284 Topic 1

Which of the following OSPF Link State Advertisements (LSA's) were created for IPV6 and do not apply to IPv4 OSPF networks? (Choose two.)

A. Link LSA (Type 8)

B. Summary LSA (Type 3)

C. Router LSA (Type 2)

D. Intra-Area Prefix LSA (Type 9)

E. Opaque LSA (Type 9)

Correct Answer: *AD*

Community vote distribution

AD (100%)

inteldarvid 5 months, 1 week ago

Selected Answer: AD

correct:
https://techhub.hpe.com/eginfolib/networking/docs/switches/5700/5998-5589r_l3-ip-rtng_cg/content/446940477.htm

upvoted 1 times

Xerath 9 months, 4 weeks ago

Selected Answer: AD

A & D are correct.

upvoted 3 times

TAZZER 11 months ago

Type 8 and type 9: Used in OSPFv3 for link-local addresses and intra-area prefixes
Correct A & D

upvoted 2 times

Noproblem22 1 year ago

A and D are correct

upvoted 2 times

Question #285 Topic 1

Router R1 has been configured with a default route like this:

```
R1#(config) ip route 0.0.0.0 0.0.0.0 10.2.3.1
```

You want to redistribute this route into OSPF but when you configure the redistribute static command under the OSPF process the default route is not present. What will create a default route in the OSPF routing process?

A. Use the redistribute static subnets command.

B. Create a default metric for the static default route.

C. Use the default-information originate command under the OSPF process.

D. Change the static default route to use an Administrative Distance (AD) greater than 110.

Correct Answer: *C*

Community vote distribution

C (100%)

HungarianDish 7 months, 2 weeks ago

Selected Answer: C

Agree with the answer

upvoted 3 times

Question #286 *Topic 1*

Which one of the following statements regarding Bidirectional Forwarding Detection (BFD) is correct?

A. BFD echo mode is the default mode of operation.

B. BFD is not supported for HSRP.

C. CEF must be disabled for BFD to work.

D. BFD is not supported when using static routes.

Correct Answer: *A*

Community vote distribution

A (100%)

inteldarvid 5 months, 1 week ago

Selected Answer: A

correct A:

https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fs_bfd.html#wp1236440

upvoted 2 times

GReddy2323 9 months, 2 weeks ago

How to Configure Bidirectional Forwarding Detection
You start a BFD process by configuring BFD on the interface. When the BFD process is started, no entries are created in the adjacency database, in other words, no BFD control packets are sent or received. BFD echo mode, which is supported in BFD Version 1 for Cisco IOS 12.4(9)T, is enabled by default. BFD echo packets are sent and received in addition to BFD control packets. The adjacency creation takes places once you have configured BFD support for the applicable routing protocols. This section contains the following procedures:
https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fs_bfd.html

upvoted 2 times

TAZZER 11 months, 1 week ago

Selected Answer: A

BFD has the following operational modes.
In asynchronous mode routers periodically send control packets to activate and maintain BFD sessions.
Asynchronous mode is available in two versions:
Asynchronous mode without echo
Asynchronous mode with echo (default value on Cisco routers)

upvoted 3 times

Question #287 Topic 1

Which of the following OSPF neighbor adjacency states is applicable only to manually configured OSPF neighbors in a Non Broadcast Multi-Access network?

A. Attempt

B. Init

C. 2-Way

D. Exstart

E. Exchange

Correct Answer: *A*

Community vote distribution

A (100%)

inteldarvid 5 months, 1 week ago

Selected Answer: A

A correct:
This state is only valid for manually configured neighbors in an NBMA environment. In Attempt state, the router sends unicast hello packets every poll interval to the neighbor, from which hellos have not been received within the dead interval.

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13685-13.html#anc14

upvoted 2 times

SujanSikrikar 9 months, 4 weeks ago

https://www.computernetworkingnotes.com/ccna-study-guide/ospf-neighbor-states-explained-with-example.html#:~:text=In%20Non-broadcast%20multi-access%20environment%20such%20as%20Frame%20Relay,it%20does%20not%20have%20to%20discover%20them%20dynamically.
Attempt

In Non-broadcast multi-access environment such as Frame Relay and X.25, OSPF uses Attempt state instead of Init state. OSPF uses this state only if neighbors are statically configured with neighbor command. In this situation, it does not have to discover them dynamically. As it already knows the neighbors, it will use unicast instead of multicast in this state.

Once neighborship is built, OSPF uses hello packets as keep alive. If a router does not receive a hello packet from any particular neighbor in dead interval, it will change its state to down from full. After changing the state it will make an effort to contact the neighbor by sending Hello packets. This effort is made in Attempt state.

Basically Both Init and Attempt states describe similar situation where one router has sent a hello packet and waiting for response.

upvoted 2 times

Question #288 *Topic 1*

With Internal BGP, there is a requirement for all peers to be logically fully meshed, where all IBGP routers must peer with all other IBGP routers. For scaling purposes, there are two mechanisms that were developed to bypass this requirement. What are they? (Choose two.)

A. Confederations

B. IBGP to EBGP route redistribution

C. BGP peer filtering

D. Route reflectors.

Correct Answer: *AD*

Community vote distribution

AD (100%)

Titini Highly Voted 👍 10 months, 1 week ago

Selected Answer: AD

Confederations: Confederations are used to break a large iBGP domain into multiple smaller sub-autonomous systems (ASs), each with its own iBGP full mesh. This reduces the number of iBGP peers in each sub-AS and allows for better scaling of the iBGP network.

Route Reflectors: Route Reflectors provide an alternative to the full iBGP mesh requirement by allowing iBGP speakers to form a partial mesh instead of a full mesh. Route Reflectors serve as a central point for iBGP speakers to exchange routing information, and iBGP speakers only need to peer with the Route Reflector instead of forming a full mesh with all iBGP speakers in the network. This reduces the number of iBGP peers and simplifies the iBGP configuration.

upvoted 5 times

mitosenoriko Most Recent 🕘 11 months, 2 weeks ago

A and D OK

upvoted 2 times

Dejjie 10 months, 1 week ago

How is A and D ok? what are your reasons?

upvoted 1 times

Question #289 *Topic 1*

You have configured policy-based routing on router R1 to force some traffic to go over an alternate link. In order to verify the configuration, which debug command should be used to verify that the specific traffic is taking the intended path?

A. Debug policy routing

B. Debug ip routing

C. Debug ip policy

D. Debug policy map

Correct Answer: *C*

Community vote distribution

C (100%)

Titini 10 months, 1 week ago

Selected Answer: C

The "debug ip policy" command is used to display PBR-related information, such as the packets matched by a PBR access-list, and the routing decisions made based on the PBR policy. This command shows whether the traffic is being correctly matched by the PBR policy, and whether it is being forwarded out the correct interface.

upvoted 4 times

Question #290 *Topic 1*

Which BGP attribute can be used to influence the path that incoming traffic takes into your AS from other Autonomous Systems?

A. Metric manipulation

B. AS_Path

C. Weight

D. Local Preference

Correct Answer: *B*

Community vote distribution

B (100%)

SAMAKEMM 2 months, 2 weeks ago

Selected Answer: B

Given answer is correct

upvoted 1 times

[Removed] 4 months, 4 weeks ago

Selected Answer: B

B
local preference and weight are locally relevant only to manipulate outgoing traffic
I can't remember much about MED

upvoted 4 times

forccnp 9 months, 1 week ago

Selected Answer: B

AS-PATH 100%

upvoted 1 times

Slinky 9 months, 3 weeks ago

Metric aka MED can also influence ingress traffic in to your autonomous system though? I agree AS_PATH is the best, but technically MED would also be correct in some situations.

upvoted 1 times

Slinky 7 months, 4 weeks ago

Commenting on my own comment here after looking through this again. The key here is autonomous-systems, plural. MED can influence ingresss if there is a single AS that has multiple ingress points to your AS. Otherwise AS_PATH or advertising a longer-prefix are your options.

upvoted 5 times

ellen_AA 11 months ago

Given answer is correct, AS-PATH to influence incoming traffic and LOCAL_PREF to influence outgoing traffic.

upvoted 4 times

Question #291

Topic 1

Routers R1 and R2 have been configured to use Bidirectional Forwarding Detection? What is the advantage of doing this?

A. It is able to discover local link failures at layer 1 and provide automatic re-routing

B. It is able to discover local link failures at layer 1 and provide automatic re-routing

C. It is able to discover local link failures at layer 1 only and provides detection for this in less than one second.

D. It is able to discover local link failures at layers 1 and 2 and provides detection for this in less than one second.

Correct Answer: *D*

Community vote distribution

D (100%)

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: D

yesy, option D:
https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fs_bfd.html

upvoted 1 times

⊟ **Titini** 10 months, 1 week ago

Selected Answer: D

D. BFD is able to discover local link failures at layers 1 and 2 and provides detection for this in less than one second. BFD can detect link failures at both the physical layer (Layer 1) and the data link layer (Layer 2) by sending and receiving packets at a very high rate (often every few milliseconds). If a BFD session detects a link failure, it can immediately signal this to the adjacent node, which can then take appropriate action to reroute traffic or take other measures to restore network connectivity.

upvoted 2 times

Question #292 Topic 1

Which BGP attribute can be used to influence the path that outgoing traffic takes from your AS to other Autonomous Systems? (Choose two.)

A. MED

B. AS_Path

C. Weight

D. Local Preference

Correct Answer: *CD*

Community vote distribution

CD (67%) BD (33%)

Titini Highly Voted 10 months, 1 week ago

Selected Answer: BD

The BGP attributes that can be used to influence the path that outgoing traffic takes from your AS to other Autonomous Systems are AS_Path and Local Preference.
The Weight attribute is useful for influencing the path of traffic within a single router, it cannot be used to influence the path of outgoing traffic from your AS to other Autonomous Systems.

upvoted 5 times

DeWalt95 Most Recent 4 weeks ago

Think B,C and D are all valid.

Local preference is shared between router in the same OS whereas Weight and AS path manipulation need to be done per router via route maps/neighbor settings.

upvoted 1 times

fizzer 3 months, 1 week ago

CD

Practically B,C and D will do the job, however, using the BGP path selection order, I would go for C and D

Priority Attribute
1 Weight
2 Local Preference
3 Originate
4 AS path length
5 Origin code

upvoted 1 times

alex711 3 months, 3 weeks ago

Selected Answer: CD

C, D is correct.

https://networklessons.com/bgp/bgp-attributes-and-path-selection

upvoted 2 times

guy276465281819372 5 months ago

Selected Answer: CD

C and D easily

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: CD

C, D is correct. AS-PATH with the prepend I can use polcy traffi inside AS

upvoted 1 times

zhlzjz 7 months ago

per long time reserch, correct anwer is B D
Weight is only right when in some special situation.
You can read offical book P519-521, Then you can find the answer.

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

I am still not sure whether it's B or C. OCG for ENARSI says:
"Weight can be set for specific routes with an inbound route map or for all routes learned from a specific neighbor. Weight is not advertised

to peers and only influences outbound traffic from a router or an AS." Then there is a picture with a topology where weight is set on two routers of the same AS, and outgoing path is only manipulated by weight from that AS. Which part are you referring to in OCG?

upvoted 2 times

HungarianDish 6 months, 3 weeks ago

I voted C, D, because weight (C) can be used for dual-home network scenarios. Weight is only locally relevant, however, if there is only one border router for eBGP peering it can be used to manipulate outbound traffic out of one AS towards other ASs.
https://community.cisco.com/t5/other-network-architecture-subjects/bgp-inbound-and-outbound-traffic/td-p/337728
I do not doubt that B (AS-path prepend) is correct, too. I just can't choose between B and C. As-path prepend used outbound is also possible.
Example:
https://blog.ipspace.net/2009/03/as-path-prepending-technical-details.html

upvoted 2 times

HungarianDish 7 months, 2 weeks ago

Selected Answer: CD

https://community.cisco.com/t5/routing/bgp-weight-local-preference-attributes-question/td-p/738421
https://community.cisco.com/t5/networking-knowledge-base/understanding-bgp-best-path-selection-manipulation/ta-p/3150576

upvoted 1 times

Almylle 5 months, 4 weeks ago

Local preference is Never Shared Between eBGP Peers, Any BGP router that receives a LOCAL_PREF attribute from an eBGP peer must ignore it (except in the case of BGP confederations)

upvoted 1 times

inteldarvid 5 months, 1 week ago

you are wrong my friend. We have policy outside our AS, with we can use LOCAL PREFERENCE for that

upvoted 1 times

drxz 7 months, 3 weeks ago

Selected Answer: CD

C and D are correct. ASpath is only for manipulating incoming traffic to your AS, not outgoing.

upvoted 2 times

HungarianDish 7 months, 2 weeks ago

However not a typical case, as-path prepend used outbound is also possible.
https://blog.ipspace.net/2009/03/as-path-prepending-technical-details.html

upvoted 2 times

Slinky 7 months, 4 weeks ago

Selected Answer: CD

C and D are correct. Weight is locally significant, but if you have one router that is dual-homed, you are still influencing egress traffic. Local pref is used for influencing egress traffic from your entire autonomous system, I.e you have two separate routers running iBGP with an uplink to a different ISP each.

upvoted 2 times

Typovy 8 months ago

Selected Answer: CD

Provided answer is correct, its Weigth and Local preference. AS Path is used to influence incoming traffic not outgoing. You can use it to influence incoming traffic but is not recomended.
@Titini "The Weight attribute is useful for influencing the path of traffic within a single router, it cannot be used to influence the path of outgoing traffic from your AS " - you dont understand what tha 'AS' is, you can have huge autonomous system network and olny one BGP router. So still you cant use Weigth?

upvoted 3 times

Typovy 8 months ago

Provided answer is correct, its Weigth and Local preference. AS Path is used to influence incoming traffic not outgoing. You can use it to influence outgoing* traffic but is not recomended.

upvoted 1 times

forccnp 9 months ago

Selected Answer: BD

B and D are correct answer

upvoted 1 times

Question #293 *Topic 1*

In order to connect between disparate OSPF and EIGRP portions of a network, mutual route redistribution has been configured. What are two disadvantages of doing this? (Choose two.)

A. Prone to routing loops if not done correctly

B. Differing metrics between the routing protocols could result in sub-optimal routing.

C. Route redistribution is not supported on most Cisco router platforms.

D. Increased convergence times.

Correct Answer: *AB*

Community vote distribution

AB (100%)

Zizu007 11 months, 2 weeks ago

Selected Answer: AB

Correct!

https://steemit.com/routing/@evaldas/the-strengths-and-weaknesses-of-redistributing-between-two-different-interior-routing-protocols

upvoted 2 times

Question #294 *Topic 1*

Which of the following statements are true regarding two EIGRP routers to become neighbors?

A. Must have identical hello and dead timers

B. Must utilize unique router ID's

C. Must have matching MTU's on the physical network links that connect the routers.

D. Must use the same ASN.

Correct Answer: *D*

Slinky 9 months, 3 weeks ago

EIGRP adjacency requires matching ASN and K-Values.

upvoted 1 times

Question #295 *Topic 1*

Which of the following statements are true regarding two OSPF routers to become neighbors? (Choose two.)

A. Must use the same ASN

B. Must have identical hello and dead timers

C. Must have matching MTU's on the physical network links that connect the routers.

D. Need not be on the same subnet.

Correct Answer: *BC*

⊟ **HarwinderSekhon** 4 months, 1 week ago

just know that there is option of "MTU ignore". But given ans is correct.

upvoted 2 times

Question #296 *Topic 1*

A new site has been added to an OPSF network using area 2. Area 2 is connected only to area 1 of this OSPF network. Area 1 is used to connect area 1 to the backbone area 0. Should you expect full connectivity to the networks located in area 2 from area 0 in this scenario?

A. Yes, by default there will be full connectivity.

B. No, you will need to redistribute the area 2 routes into area 0.

C. No, a virtual link is needed to logically connect area 2 info area 0.

D. Yes, but area 2 will need to be configured as a stub area.

Correct Answer: *C*

Community vote distribution

C (100%)

⊟ **spada05** 5 months ago

Selected Answer: C

Typo? I assume it meant to say "Area 2 is used to connect area 1..." instead of area 1 to area 1. Given that, a virtual link is correct.

upvoted 2 times

Question #297

Topic 1

Router R1 is configured using VRF's to support customer VPN's. Some customers are using the same private IP address space. Which of the following is used to ensure that these routes are unique when advertised throughout the VPN?

A. Route Distinguisher

B. Route Targets

C. MP-BGP

D. LDP

Correct Answer: *A*

Community vote distribution

A (100%)

Malasxd 7 months, 1 week ago

Selected Answer: A

A for sure

upvoted 2 times

Question #298

Topic 1

Two MPLS routers, R1 and R2, are not directly connected and have an established LDP session running between them. What type of LDP session is this?

A. Remote LDP session

B. Direct LDP session

C. Tunneled LDP session

D. Targeted LDP session

Correct Answer: *D*

Community vote distribution

D (100%)

chris110 3 months, 2 weeks ago

Selected Answer: D

In the context of MPLS (Multiprotocol Label Switching), when two routers that are not directly connected have an established LDP (Label Distribution Protocol) session, it is referred to as a "Targeted LDP session."

So, the correct answer is:

D. Targeted LDP session

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: D

yes!!!! option D

upvoted 2 times

Wooker 8 months ago

Selected Answer: D

https://www.cisco.com/en/US/docs/ios-xml/ios/mp_ldp/configuration/15-2s/mp-ldp-overview.html

upvoted 2 times

Question #299 *Topic 1*

What is the total length of an MPLS header?

A. 16 bits

B. 20 bits

C. 28 bits

D. 32 bits

Correct Answer: *D*

Community vote distribution

D (100%)

⊟ **Titini** 10 months, 1 week ago

Selected Answer: D

The total length of an MPLS header is 32 bits, or 4 bytes. The MPLS header is composed of several fields, including a label, experimental bits, a bottom-of-stack (BoS) bit, a time-to-live (TTL) field, and a type field. Each field is a specific length, and together they comprise the 32-bit MPLS header.

upvoted 3 times

Question #300 *Topic 1*

Which of the following are valid fields in an MPLS header? (Choose four.)

A. Label

B. Sequence Number

C. Experimental (Exp)

D. Bottom of Stack (BoS)

E. Time to Live (TTL)

F. Checksum

Correct Answer: *A, C, D, E*

⊟ **inteldarvid** 5 months, 1 week ago

yesy corerct
Label value: the name says it all, this is where you will find the value of the label.
EXP: these are the three experimental bits. These are used for QoS, normally the IP precedence value of the IP packet will be copied here.
S: this is the “bottom of stack” bit. With MPLS it’s possible to add more than one label, you’ll see why in some of the MPLS VPN lessons. When this bit is set to one, it’s the last MPLS header. When it’s set to zero then there is one or more MPLS headers left.
TTL: just like in the IP header, this is the time to live field. You can use this for traces in the MPLS network. Each hop decrements the TTL by one.

upvoted 3 times

Question #301 *Topic 1*

What TCP port is used by LDP to provide for reliable transport connections?

A. 646

B. 648

C. 752

D. 712

Correct Answer: *A*

Noproblem22 1 year ago

A is the correct answer

upvoted 4 times

Question #302 *Topic 1*

Which of the following are control plane protocols used within a service provider MPLS network? (Choose two.)

A. OAM

B. RSVP

C. Targeted LDP

D. SNMP

E. LDP

Correct Answer: *BE*

HarwinderSekhon 4 months, 1 week ago

Control plane protocols are responsible for the operations and signaling within a network, including in an MPLS (Multiprotocol Label Switching) environment.

The correct answers from the given options are:

B. RSVP (Resource Reservation Protocol)
E. LDP (Label Distribution Protocol)

Explanation:

RSVP: It's used for reserving resources and is also employed as a signaling protocol in MPLS Traffic Engineering (TE).
LDP: This protocol is used for label distribution in MPLS, enabling routers to learn the labels to use for forwarding traffic.
Option C (Targeted LDP) is a specialized use of LDP, so it's related but not a distinct protocol in this context.

Options A (OAM - Operations, Administration, and Maintenance) and D (SNMP - Simple Network Management Protocol) are not typically classified as control plane protocols within the context of an MPLS network, as they are more related to management and monitoring.

upvoted 1 times

DUBC89x 1 year ago

Given anser is correct.
There are two standardized protocols for managing MPLS paths: the Label Distribution Protocol (LDP) and RSVP-TE, an extension of the Resource Reservation Protocol (RSVP) for traffic engineering. Furthermore, there exist extensions of the Border Gateway Protocol (BGP) that can be used to manage an MPLS path.

upvoted 1 times

Question #303 *Topic 1*

In a typical MPLS VPN, which routers act as the MPLS label imposition and disposition points in the network?

A. CE Router

B. P router

C. PE Router

D. Core router

Correct Answer: *C*

Community vote distribution

C (100%)

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: C

yes correct is C:
Provider Edge Routers (PE) operate at the edge of the provider network. They perform Label Edge Router (LER) imposition and disposition operations at the edge of an MPLS network. In an MPLS network, the ingress edge router receives the packet and adds a label to the packet. The egress edge router removes the label.

https://content.cisco.com/chapter.sjs?uri=%2Fsearchable%2Fchapter%2Fwww.cisco.com%2Fcontent%2Fen%2Fus%2Ftd%2Fdocs%2Fswitches%2Fmetro%2Fme3600x_3800x%2Fsoftware%2Frelease%2F15-1_2_ey%2Fconfiguration%2Fguide%2F3800x3600xscg%2Fswmpls.html.xml#:~:text=They%20perform%20Label%20Edge%20Router,edge%20router%20removes%20the%20label.

upvoted 2 times

Question #304 *Topic 1*

In an MPLS network, which of the following describes the role of the Provider (P) router?

A. To connect to customer edge (CE) devices

B. To connect to PE routers and act as transit routers

C. To impose MPLS labels

D. To filter VPN routes in the core

Correct Answer: *B*

Community vote distribution

B (100%)

⊟ **inteldarvid** 5 months ago

Selected Answer: B

B correct

upvoted 2 times

Question #305 *Topic 1*

The MPLS LDP autoconfiguration feature allows you to enable LDP on every interface that is associated with an IGP instance. Which of the following Interior Gateway Protocols support this? (Choose two.)

A. OSPF

B. IS-IS

C. BGP

D. RIP

E. EIGRP

Correct Answer: *AB*

alex711 3 months, 3 weeks ago

The given answer is correct

upvoted 1 times

GReddy2323 9 months, 2 weeks ago

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/15-s/mp-ldp-15-s-book/mp-ldp-autoconfig.pdf

upvoted 1 times

Question #306 *Topic 1*

In an MPLS VPN network, how are customer routes controlled and distributed?

A. Through the use of GRE tunnels

B. Customer routes are redistributed into the IGP that the service provider is using

C. Customer routes are redistributed into BGP within the service provider

D. It is distributed through the use of route targets

Correct Answer: *D*

Community vote distribution

D (67%) C (33%)

fizzer 3 months, 1 week ago

Selected Answer: D

Correct answer is D (Route Target)

High level steps involved in sharing customer route from Site A to Site B

Site A PE learns the routes from CE via IGP
They are redistributed into BGP under the context of customer vrf
To share these routes across to Ps and other PEs, they are exported into Global BGP table "route-target export 11:22"
Now in Global BGP table on Site B's PE, they are imported into the customer vrf table with command "route-target import 11:22"
They are then redistributed into the IGP between the vrf and Site B CE
Same process to share Site B routes to A but with B setting different numeric value for the RT

Yes, redistribution (technical word) between BGP and the IGPs is part of the process, but MP-BGP "controls and distributes" (non technical words) customer routes using RT

upvoted 1 times

alex711 3 months, 3 weeks ago

Selected Answer: D

It is D.

https://notes.networklessons.com/mpls-route-target

upvoted 2 times

HarwinderSekhon 4 months, 1 week ago

Selected Answer: C

I was going for D until I found this in official cisco doc
•Extended MP-BGP community attributes are used to control the distribution of customer routes.

https://content.cisco.com/chapter.sjs?uri=/searchable/chapter/www.cisco.com/content/en/us/td/docs/net_mgmt/prime/fulfillment/6-2/theory/operations/guide/theory/mpls.html.xml#:~:text=Extended%20MP%2DBGP%20community%20attributes,the%20distribution%20of%20customer%20routes.&text=Each%20customer%20route%20is%20associated,correct%20egress%20customer%20edge%20router.

upvoted 1 times

fizzer 3 months, 1 week ago

Route Target is an extended community attribute in MP-BGP

upvoted 1 times

[Removed] 4 months, 4 weeks ago

I feel like this should be a choose two kind of question, D is t the only part of customer Route distribution, MPIBGP uses the vpn4 address to redistribute the addresses.

upvoted 2 times

[Removed] 4 months, 4 weeks ago

I had a typo, "D is not the only part of control and distribution"

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

Selected Answer: D

"D" is more accurate.
https://notes.networklessons.com/mpls-route-target
MPLS Route Target
When implementing MPLS VPNs, we use a Route Target or RT to allow PE routers to control the distribution of VPN routes to the appropriate VRFs.
Now an RT is actually a BGP extended community attribute that is used to control the distribution of VPN routing information between PE routers.

upvoted 3 times

HungarianDish 6 months, 2 weeks ago

Similar to Question #117 .
https://www.ccexpert.us/mpls-design/chapter-5-packetbased-mpls-vpns.html
The distribution of VPN routing information is controlled through the use of VPN route target communities...

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: C

https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/BRKMPL-1100.pdf
• Between PE routers: customer routes exchanged via BGP (MP-BGP)

https://www.rfc-editor.org/rfc/rfc2547.txt
"PE routers use BGP to distribute VPN routes to each other"
"the "Route Target" attribute ... identifies only a set of sites which will be able to use the route"

https://networklessons.com/mpls/mpls-layer-3-vpn-explained
The PE router will then redistribute everything in BGP (Multi Protocol BGP).

https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/BRKMPL-1100.pdf
• MP-iBGP: advertises VPNv4 prefixes + labels
"Route Targets dictate which VRF will receive what routes"

Interactions Between VRF and BGP VPN Signaling:
PE1 redistributes VPNv4 route into MP-iBGP

upvoted 2 times

robi1020 7 months, 1 week ago

The Q is how are customer routes controlled and distributed?

D is correct. This is a classic cisco Q... I feel you tho

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Hi Robi, I can accept route targets (D) as the answer because of the word "controlled". Still, MP-BGP could be an answer to how are routes distributed.

upvoted 1 times

Question #307 *Topic 1*

At which layer of the OSI model is an MPLS label imposed?

A. Layer 2

B. Layer 3

C. Between layers 2 and 3

D. Between layers 3 and 4

Correct Answer: *C*

Community vote distribution

C (100%)

[Removed] 4 months, 4 weeks ago

Selected Answer: C

The label is a shim header between L2 and L3

upvoted 2 times

Noproblem22 1 year ago

C is right

upvoted 1 times

Question #308 *Topic 1*

Which of the following are valid IPv6 Router Advertisement (RA) Guard modes? (Choose two.)

A. Guard mode

B. Host mode

C. Router mode

D. Open mode

E. Closed mode

Correct Answer: *BC*

Community vote distribution

BC (100%)

inteldarvid 4 months, 3 weeks ago

Selected Answer: BC

correct B nd C

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/15-mt/ip6f-15-mt-book/ip6-ra-guard.html

upvoted 1 times

Malasxd 7 months, 2 weeks ago

Selected Answer: BC

B, C
There are two primary RA Guard modes:

Host mode (Default mode): In this mode, the RA Guard function expects that trusted ports are connected to hosts that do not send any Router Advertisement or Router Redirect messages. This mode will drop all such messages if received on a trusted port.

Router mode: In this mode, the RA Guard function allows trusted ports to send Router Advertisement and Router Redirect messages. This mode is configured if the trusted port is connected to a legitimate router.

upvoted 3 times

Question #309 *Topic 1*

Which of the following statements are true regarding the e IPv6 RA Guard feature?

A. This feature is support on LAG bundles interfaces

B. This feature is supported on private VLANs

C. Packets dropped by the IPv6 RA Guard feature cannot be spanned.

D. This feature offers protection in networks where IPv6 traffic is tunneled.

Correct Answer: *B*

Community vote distribution

B (100%)

DUBC89x 1 year ago

Selected Answer: B

Given Answer is correct.
"https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-16/ip6f-xe-16-book/ip6-ra-guard.html"

upvoted 4 times

Question #310 Topic 1

Unicast Reverse Path Forwarding (uRPF) has been configured on a service provider network to protect itself from spoofed based attacks. Which of the following are valid uRPF modes? (Choose two.)

A. Strict mode

B. Open mode

C. Closed mode

D. Block mode

E. Loose mode

Correct Answer: *AE*

Community vote distribution

AE (100%)

sajjad_gayyem 6 months ago

The answer is true. unicast Revers Path Forwarding is a security feature against Source IP spoofing.
Good explanation can be found here:
https://networklessons.com/cisco/ccie-routing-switching/unicast-reverse-path-forwarding-urpf

upvoted 3 times

Xerath 9 months, 3 weeks ago

Selected Answer: AE

The given answer is correct.

upvoted 2 times

Question #311 *Topic 1*

Which of the following are commonly used ports when implementing RADIUS based authentication and accounting? (Choose two.)

A. UDP port 1644 for authentication

B. UDP port 1812 for authentication

C. TCP port 1812 for authentication

D. UDP port 1813 for accounting

E. TCP port 1813 for accounting

F. UDP port 1644 for accounting

Correct Answer: *BD*

Community vote distribution

BD (100%)

alex711 3 months, 3 weeks ago

Selected Answer: BD

given answer is correct.

upvoted 1 times

Zizu007 11 months, 2 weeks ago

Selected Answer: BD

Correct!

https://community.cisco.com/t5/routing/which-port-numbers-are-used-for-radius-accounting-and-radius/td-p/2494536

upvoted 2 times

Question #312 *Topic 1*

Which of the following are valid restrictions when configuring Control Plane Policing (CoPP) on Cisco devices? (Choose two.)

A. You cannot use the “log” keyword with CoPP on the access list entries

B. CEF must be disabled

C. The only match types supported with CoPP is ip precedence, ip dscp, and access-group

D. Only standard access-lists are supported.

Correct Answer: *AC*

Community vote distribution

AC (100%)

alex711 3 months, 3 weeks ago

Selected Answer: AC

AC is OK.

upvoted 1 times

GReddy2323 6 months ago

Selected Answer: AC

Answer is correct
https://www.cisco.com/en/US/docs/general/Test/dwerblo/broken_guide/copp.html
•To avoid matching the filtering and policing that are configured in a subsequent class, configure policing in each class. CoPP does not apply the filtering in a class that does not contain a police command. A class without a police command matches no traffic.

•The ACLs used for classification are QoS ACLs. QoS ACLs supported are IP standard, extended, and named (IPv6 ACLs are not supported in hardware).

•These are the only match types supported:

–ip precedence

–ip dscp

–access-group

•CoPP does not support ACEs with the log keyword.

upvoted 1 times

sasasan12345 9 months, 1 week ago

A and D are correct.
Defining Interesting Traffic with Extended ACLs.

upvoted 1 times

[Removed] 4 months, 4 weeks ago

You mean AC, D is specifying standard lists

upvoted 1 times

mitosenoriko 11 months, 2 weeks ago

A and C correct
Dont use "log"
support type
– ip precedence
– ip dscp
– access-group

upvoted 2 times

Question #313 *Topic 1*

Which of the following are used to validate the source of IPv6 traffic and are considered IPv6 layer 2 snooping features? (Choose two.)

A. DHCPv6 Guard

B. DHCPv6 Root Guard

C. IPv6 Source Guard

D. IPv6 Prefix Guard

Correct Answer: *CD*

Community vote distribution

CD (100%)

je2004 7 months, 2 weeks ago

Agree.

https://content.cisco.com/chapter.sjs?uri=/searchable/chapter/content/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-3e/ip6f-xe-3e-book/ip6f-xe-3e-book_chapter_0110.html.xml

upvoted 2 times

Malasxd 7 months, 2 weeks ago

Selected Answer: CD

C, D seems more correct for me.

I would say DHCPv6 also prevents snooping, but the question especify "layer 2 snooping" so just Source Guard and prefix Guard work with binding table.

upvoted 2 times

Question #314 *Topic 1*

You want to implement AAA on router R1 for a more robust authentication and authorization system. What is typically the first global command used to do this?

A. aaa new-model

B. aaa enable

C. aaa server-group

D. aaa authentication login

Correct Answer: *A*

Community vote distribution

A (100%)

Brand 4 months, 1 week ago

Omg look at that question... Just amazing...

upvoted 1 times

forccnp 10 months ago

Selected Answer: A

A is correct

upvoted 2 times

Question #315 *Topic 1*

A time based access list has been configured on R1 to allow SSH access to the device only on weekdays. Which of the following are valid options when using the time range command? (Choose two.)

A. relative

B. recurring

C. absolute

D. periodic

Correct Answer: *CD*

Community vote distribution

CD (100%)

inteldarvid 5 months, 1 week ago

Selected Answer: CD

C, D correct:
https://www.cisco.com/c/en/us/support/docs/security/ios-firewall/23602-confaccesslists.html

upvoted 2 times

Malasxd 7 months, 2 weeks ago

Selected Answer: CD

C and D correct

```
router(config)#time-range teste
router(config-time-range)#?
Time range configuration commands:
absolute absolute time and date
default Set a command to its defaults
exit Exit from time-range configuration mode
no Negate a command or set its defaults
periodic periodic time and date
```

upvoted 1 times

GReddy2323 9 months, 2 weeks ago

Selected Answer: CD

https://www.cisco.com/c/en/us/support/docs/smb/switches/cisco-small-business-300-series-managed-switches/smb5660-configure-time-range-settings-on-a-switch-through-the-comman.html

upvoted 1 times

mitosenoriko 11 months, 2 weeks ago

C and D correct

```
   (config)# time-range name
   (config-time-range)# periodic days-of-the-week hh:mm to [days-of-the-week] hh:mm
   or
   (config-time-range)# absolute [ start time date ] [ end time date ]
```

upvoted 3 times

Question #316 *Topic 1*

First-Hop Security (FHS) is a set of features to optimize IPv6 link operation, and help with scale in large L2 domains. Which of the following are valid First-Hop Security features supported by Cisco? (Choose three.)

A. IPv6 RA Guard

B. IPv6 Source Guard

C. DHCPv6 Guard

D. IPv6 Snooping

E. DHCPv6 Snooping

Correct Answer: *ACD*

Community vote distribution

ACD (100%)

⊟ **Muste** 4 months, 2 weeks ago

Selected Answer: ACD

the question is asking about L2 domains that's why B isn't qualified

upvoted 4 times

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: ACD

A, C, D : Correct:
https://www.cisco.com/c/en/us/td/docs/dcn/nx-os/nexus9000/102x/configuration/Security/cisco-nexus-9000-nx-os-security-configuration-guide-102x/m-configuring-ipv6-first-hop-security.html

upvoted 1 times

⊟ **slcc99** 5 months, 3 weeks ago

I think A, B, and C are correct.Because "The IPv6 Snooping Policy feature is deprecated and the Switch Integrated Security Feature (SISF)-based device tracking feature replaces it and offers the same capabilities."

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9300/software/release/17-1/configuration_guide/sec/b_171_sec_9300_cg/configuring_ipv6_first_hop_security.html#:~:text=The%20IPv6%20Snooping%20Policy%20feature%20is%20deprecated%20and%20the%20Switch%20Integrated%20Security%20Feature%20(SISF)%2Dbased%20device%20tracking%20feature%20replaces%20it%20and%20offers%20the%20same%20capabilities.

upvoted 1 times

⊟ **HungarianDish** 7 months, 1 week ago

A,B,C,D!!!
RA Guard, DHCPv6 Guard, Source Guard, IPv6 ND snooping = device-tracking

https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/BRKSEC-3200.pdf
https://networklessons.com/cisco/ccie-routing-switching-written/ipv6-first-hop-security-features
https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/15S/configuration/guide/7600_15_0s_book/IPv6_Security.html

upvoted 2 times

⊟ **Patrick1234** 11 months ago

Source Guard is also part of the FHS features, however, it needs IPv6 Snooping to be enabled... I would not know why you should not pick that one as well, but i guess it's save to use the given answer here... So A, C, and D seem to be correct.

upvoted 1 times

⊟ **DUBC89x** 1 year ago

Selected Answer: ACD

Given answer is correct
"https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/15S/configuration/guide/7600_15_0s_book/IPv6_Security.html"

upvoted 1 times

Question #317 *Topic 1*

What are the four stages of obtaining an IP address from a DHCP server that corresponds to the acronym DORA?

A. Discover, Offer, Release, Addressing

B. Discover, Obtain, Request, Acknowledge

C. Determine, Offer, Release, Acknowledge

D. Discover, Offer, Request, Acknowledge

Correct Answer: *D*

Question #318 *Topic 1*

SNMPv2 has been used throughout a network to manage all of the network devices. You have been asked to migrate to an SNMPv3 solution instead. What is the biggest advantage to migrating from SNMPv2 to SNMPv3?

A. Enhanced security, including encryption of passwords

B. Enhanced performance, supporting more messages per minute.

C. Enhanced scaling, supporting thousands more devices per network segment than SNMPv2.

D. Using a push model instead of pull. SNMPv3 uses telemetry to push data to SNMP management stations in real time.

Correct Answer: *A*

mitosenoriko 11 months, 2 weeks ago

A is correct.
SNMPv3 available Encryption.

upvoted 1 times

Noproblem22 1 year ago

A is correct

upvoted 4 times

Question #319 *Topic 1*

You are configuring Netflow on various network elements in order to gain visibility into the traffic types used. How many export destinations can this Network data be sent to?

A. Up to 2

B. Up to 4

C. Up to 8

D. There is no limitation on the number of flow data export destinations.

Correct Answer: *A*

HungarianDish 6 months, 2 weeks ago

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/fnetflow/configuration/xe-16-10/fnf-xe-16-10-book/fnf-output-features.pdf
Original NetFlow is limited to only two export destinations per cache.

upvoted 2 times

mitosenoriko 11 months, 2 weeks ago

A is correct
max is 2

upvoted 1 times

Zizu007 11 months, 3 weeks ago

answer is correct:

**You can configure a maximum of two export destinations for NetFlow.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/netflow/configuration/15-mt/nf-15-mt-book/get-start-cfg-nflow.html

upvoted 2 times

Question #320 *Topic 1*

A Cisco router has just been configured for NTP and is synchronized with the configured NTP server. However, log messages still show an incorrect time. What else should be done to match the log messages time stamps with the NTP based time?

A. Wait a bit longer for the synchronized time to get applied to new log messages.

B. Configure the "service timestamps log datetime localtime" command in global mode.

C. Configure the "service timestamps log datetime synchronize" command globally

D. Configure the "service timestamps log ntp" command in global config mode.

Correct Answer: *B*

Zizu007 11 months, 3 weeks ago

Selected Answer: B

Correct!

https://community.cisco.com/t5/networking-knowledge-base/router-log-timestamp-entries-are-different-from-the-system-clock/ta-p/3132258

upvoted 1 times

Question #321

Topic 1

There is an issue between two nodes within your network, and you are using Cisco DNA Center Path Trace to help troubleshoot the problem. Which of the following statements are true regarding the Path Trace tool?

A. Overlapping IP addresses are supported.

B. Path trace between a fabric client and a non-fabric client is supported

C. Path trace between a wired client and a wireless client is supported

D. Only TCP traffic is supported.

Correct Answer: *C*

Malasxd 7 months, 1 week ago

Selected Answer: C

Given answer is correct.
https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center-assurance/2-2-2/b_cisco_dna_assurance_2_2_2_ug/b_cisco_dna_assurance_2_2_2_ug_chapter_01111.html

upvoted 1 times

mitosenoriko 11 months, 2 weeks ago

C is correct
Overlapping IP addresses not supported.
Path trace between a fabric client and a non-fabric client is not supported.
TCP and UDP are supported.

upvoted 3 times

Question #322 *Topic 1*

Which of the following are valid DHCP options that DHCP servers can be configured to use with DHCP clients when offering a lease? (Choose two.)

A. DHCP Option 1: subnet mask

B. DHCP Option 3: Lease Duration

C. DHCP Option 4: Client host name

D. DHCP Option 6: DNS servers

Correct Answer: *AD*

alex711 3 months, 3 weeks ago

Selected Answer: AD

The given answer is correct.

upvoted 1 times

R9_9_9 8 months, 3 weeks ago

Subnet Mask 1
Time Offset 2
Router 3
Time Server 4
Name Server Option 5
Domain Name Server 6
https://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/network_registrar/9-1/dhcp/guide/DHCP_Guide/DHCP_Guide_appendix_01101.pdf
Page 2 and 3

upvoted 4 times

sasasan12345 9 months, 1 week ago

3 is Router(Default Gate)
4 is Time Server.

upvoted 1 times

sasasan12345 9 months, 1 week ago

Collect!

upvoted 2 times

mitosenoriko 11 months, 2 weeks ago

A and C and D correct...i dont know.

upvoted 1 times

Zizu007 11 months, 3 weeks ago

Correct!

Page2
https://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/network_registrar/9-1/dhcp/guide/DHCP_Guide/DHCP_Guide_appendix_01101.pdf

upvoted 1 times

Question #323 *Topic 1*

Which feature of the Cisco DNA Center allows you to run diagnostic CLI commands to the devices that are managed by DNA Center for troubleshooting purposes?

A. Command Runner

B. DNA Spaces

C. DNA Advantage

D. Intelligent Capture

Correct Answer: *A*

⊟ **Malasxd** 7 months, 2 weeks ago

Selected Answer: A

Command runner is the feature in Cisco DNA Center that allows you to execute commands on the devices managed by Cisco DNA Center

https://developer.cisco.com/docs/dna-center/#!command-runner

upvoted 1 times

Question #324 *Topic 1*

You want to change the Administrative Distance of external EIGRP routes from the default of 170 to 130 instead on router R1 while leaving the default AD value for internal EIGRP routes. Which set of command will accomplish this?

A. R1(config)#router eigrp -

R1(config-router)#distance 170 -

B. R1(config)#router eigrp 1 -
R1(config-router)#distance eigrp 90 130

C. R1(config)#router eigrp 1 -
R1(config-router)#distance eigrp 130 90

D. R1(config)#router eigrp 1 -
R1(config-router)#distance 90 130

Correct Answer: *B*

⊟ **JJH3003** Highly Voted 4 months, 3 weeks ago

NO!

```
R1(config)#router eigrp 1
R1(config-router)#distance 90 130
^
% Invalid input detected at '^' marker.
R1(config-router)#
================================
R1(config)#router eigrp 1
R1(config-router)#distance ?
<1-255> Set route administrative distance
eigrp Set distance for internal and external routes

R1(config-router)#distance eigrp ?
<1-255> Distance for internal routes

R1(config-router)#distance eigrp 90 ?
<1-255> Distance for external routes

R1(config-router)#distance eigrp 90 130 ?
<cr>
===================================
```

Answer B is correct!

upvoted 5 times

⊟ **HarwinderSekhon** Most Recent 4 months, 1 week ago

Selected Answer: B

eigrp keyword needs to be there otherwise You will modify AD of EIGRP

upvoted 1 times

⊟ **JieW** 4 months, 2 weeks ago

Selected Answer: B

JJH3003 is correct. You need eigrp in the command otherwise it comes up as setting AD

upvoted 1 times

⊟ **[Removed]** 4 months, 3 weeks ago

Selected Answer: D

Correct answer is D. The command does not require the EIGRP keyword

Under the EIGRP routing process configure

Distance (internal AD) (external AD)

upvoted 1 times

Question #325 Topic 1

Which of the following are valid TFTP error codes? (Choose two.)

A. Error Code 1 – File not found

B. Error Code 2 – Unknown error

C. Error code 3 – Invalid user

D. Error code 6 – File already exists

E. Error code 8 – Undefined error

Correct Answer: *AD*

Jey117 2 months, 1 week ago

What does this have to do with you being a Network admin? Can't you justo quickly Google it instead of memorizing it. WTF Cisco

upvoted 4 times

HarwinderSekhon 4 months, 1 week ago

0 Not defined, see error message (if any).
1 File not found.
2 Access violation.
3 Disk full or allocation exceeded.
4 Illegal TFTP operation.
5 Unknown transfer ID.
6 File already exists.
7 No such user.

upvoted 3 times

kaisehhop 11 months ago

Correct!
https://www.cisco.com/c/en/us/support/docs/ip/trivial-file-transfer-protocol-tftp/13705-15.html

upvoted 2 times

Question #326 *Topic 1*

What are the two prerequisites of setting up DMVPN tunnel? (Choose two.)

A. Before a multipoint GRE (mGRE) and IPsec tunnel can be established, define an Internet Key Exchange (IKE) policy by using the crypto isakmp policy command.

B. The Public IP's of the routers should be able to ping each other.

C. To enable 2547oDMPVN - Traffic Segmentation Within DMVPN configure multiprotocol label switching (MPLS) by using the mpls ip command

D. It is mandatory to use wildcard preshared keys to build the DMVPN tunnel

E. DMVPN can work on all OEM devices that support IKE.

Correct Answer: *AC*

HungarianDish 6 months, 2 weeks ago

The question is clearly taken from here, as DUBC89x pointed out, and so, I agree on the answers "A", "C".
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_dmvpn/configuration/15-mt/sec-conn-dmvpn-15-mt-book/sec-conn-dmvpn-dmvpn.html#GUID-D8F6839F-D735-4C8E-A199-602CDD8F7DD0

However: IPsec is only optional for basic DMVPN tunnel configuration.
https://networklessons.com/cisco/ccie-routing-switching/dmvpn-over-ipsec

Of course, I can't imagine using DMVPN without IPsec, still it is a tricky question, because IPsec is not needed for the DMVPN tunnel establishment. Also, normally I would check reachability via the WAN/public IPs before setting up the tunnel. So, I would not say that "B" is wrong.

upvoted 1 times

alex711 3 months, 3 weeks ago

Yes, Agree.

upvoted 1 times

DUBC89x 1 year ago

Given answer is correct.
"Prerequisites for Dynamic Multipoint VPN (DMVPN)
Before a multipoint GRE (mGRE) and IPsec tunnel can be established, you must define an Internet Key Exchange (IKE) policy by using the crypto isakmp policy command.

For the NAT-Transparency Aware enhancement to work, you must use IPsec transport mode on the transform set. Also, even though NAT-Transparency can support two peers (IKE and IPsec) being translated to the same IP address (using the User Datagram Protocol [UDP] ports to differentiate them [that is, Peer Address Translation (PAT)]), this functionality is not supported for DMVPN. All DMVPN spokes must have a unique IP address after they have been NAT translated. They can have the same IP address before they are NAT translated.

To enable 2547oDMPVN--Traffic Segmentation Within DMVPN you must configure multiprotocol label switching (MPLS) by using the mpls ip command."
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_dmvpn/configuration/15-mt/sec-conn-dmvpn-15-mt-book/sec-conn-dmvpn-dmvpn.html#GUID-D8F6839F-D735-4C8E-A199-602CDD8F7DD0

upvoted 1 times

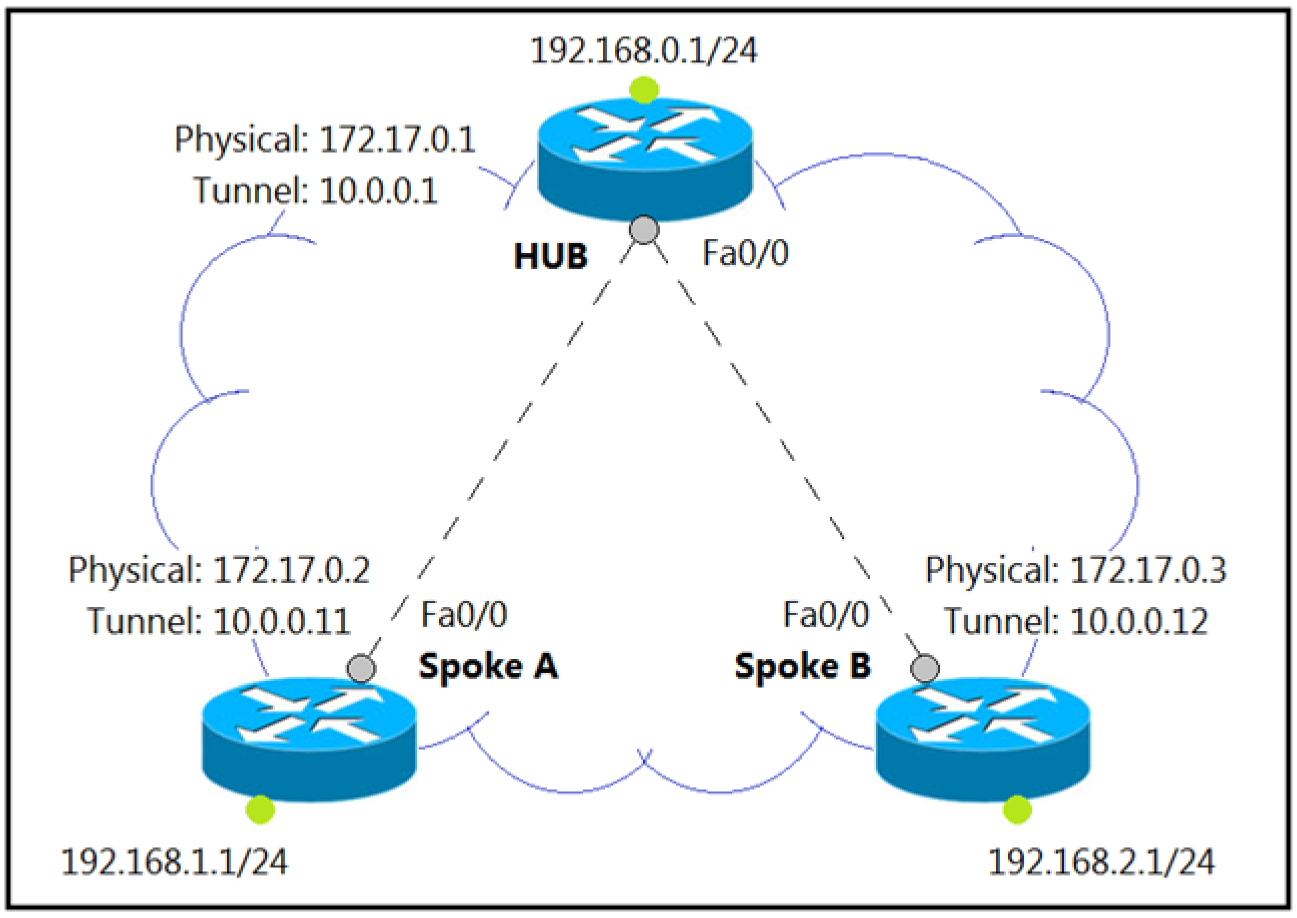

Refer to the exhibit. An administrator is setting up above shown routers to enable MVPN with mGRE mode. What would be the recommended interface configuration that must be done by the engineer to make it to work?

A. interface Tunnel0
description mGRE - DMVPN Tunnel
ip address 10.0.0.1 255.255.255.0
ip nhrp map multicast dynamic
ip nhrp network-id 1
tunnel source 10.0.0.1
tunnel mode IPSec multipoint

B. interface Tunnel0
description mGRE - DMVPN Tunnel
ip address 10.0.0.1 255.255.255.0
ip nhrp map multicast dynamic
ip nhrp network-id 1
tunnel source 10.0.0.1
tunnel mode gre multipoint

C. interface Tunnel0
description mGRE - DMVPN Tunnel
ip address 10.0.0.1 255.255.255.0
ip nhrp network-id 1
tunnel source 172.17.0.1
tunnel mode IPsec multipoint

D. interface Tunnel0
description mGRE - DMVPN Tunnel
ip address 10.0.0.1 255.255.255.0
ip nhrp map multicast dynamic
ip nhrp network-id 1
tunnel source 10.0.0.1
tunnel destination 172.17.0.2
tunnel mode IPsec multipoint

Correct Answer: *B*

kebkim Highly Voted 👍 1 year ago

Answer is C. The same question is 113.

upvoted 7 times

raw007 6 months, 3 weeks ago

But C is configured with GRE Multipoint

upvoted 1 times

kaupz Most Recent 🕒 1 week, 2 days ago

there is no right answer...

ACD are excluded because:

```
R1(config-if)#tunnel mode ipsec multipoint
^
% Invalid input detected at '^' marker.
```

And B is excluded because tunnel source cannot be tunnel itself

upvoted 1 times

Brand 3 months, 2 weeks ago

"tunnel mode IPsec multipoint" doesn't seem like a valid command at all. But the options A and B using tunnel source as tunnel IP itself so they can't be correct too. WTF is this nonsense...

```
C1-HUB(config-if)#tunnel mode ipsec ?
ipv4 over IPv4
ipv6 over IPv6

C1-HUB(config-if)#tunnel mode ipsec ipv4 ?
v6-overlay Overlay traffic v6
<cr>

C1-HUB(config-if)#tunnel mode ipsec ipv4
```

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

100 % is C because 127.17.0.1 is a NBMA PUBLIC soruce

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: C

"tunnel source" = physical interface

tunnel mode gre multipoint
Do one of the following:
tunnel protection ipsec profile name
tunnel protection psk key

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_dmvpn/configuration/xe-16-11/sec-conn-dmvpn-xe-16-11-book/sec-conn-dmvpn-dmvpn.html

upvoted 1 times

forccnp 9 months, 1 week ago

Selected Answer: C

Tunel source should be physical address

upvoted 1 times

Lilienen 10 months, 2 weeks ago

Selected Answer: C

Answer is C

upvoted 1 times

sylvesterbello1 11 months, 3 weeks ago

C is the correct answer.
Physical interface ip is the source interface, not the tunnel ip

upvoted 2 times

jarz 1 year ago

Selected Answer: C

Answer is C

upvoted 3 times

DUBC89x 1 year ago

Answer is C
You cannot have a tunnel destination from the Tunnel IP address, has to be the address of the physical link.

upvoted 3 times

Question #328

Topic 1

Select three benefits of setting up a MPLS Network from the below options. (Choose three.)

A. Connection less Service

B. Security as good as connection-oriented VPNs

C. Provides IPS level intelligence to filter packets.

D. Integrated QoS support

E. All variations of Static routes are supported

Correct Answer: *ABD*

mitosenoriko 11 months, 2 weeks ago

ABD is correct

upvoted 2 times

Question #329 *Topic 1*

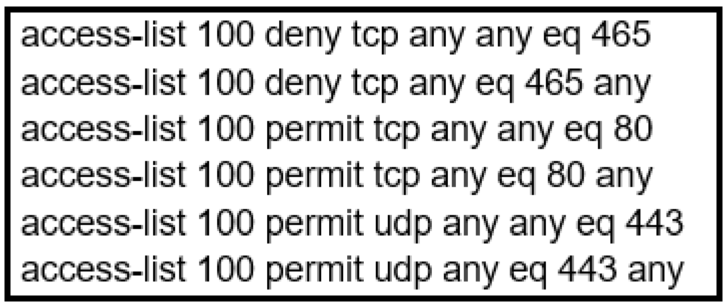

```
access-list 100 deny tcp any any eq 465
access-list 100 deny tcp any eq 465 any
access-list 100 permit tcp any any eq 80
access-list 100 permit tcp any eq 80 any
access-list 100 permit udp any any eq 443
access-list 100 permit udp any eq 443 any
```

Refer to the Exhibit. The access-lists are configured on the network device. There is a server behind the network device. User are trying to access the server securely however they are not able to access it. What changes would you recommend to the above configuration?

A. Permit tcp port 465

B. Permit tcp port 3389

C. Permit tcp port 443

D. Permit tcp any any

Correct Answer: *C*

forccnp 10 months ago

given answer is correct.
same question #125

upvoted 2 times

ellen_AA 11 months, 2 weeks ago

Permit tcp port 443 doesn't exist!!
should be: permit tcp ip any any eq 443

upvoted 2 times

Question #330 *Topic 1*

Which of the following is true regarding IPsec Pre-fragmentation (Look-Ahead Fragmentation)? (Choose two.)

A. Operates in tunnel mode only

B. Operates in transport mode only

C. Is used to help in the overall IPsec throughput since the end host is able to avoid packet reassembly after packet decryption.

D. Is not dependent on the MTU of the physical interface used for IPsec.

E. Does not support Path MTU Discovery

Correct Answer: *AC*

alex711 3 months, 3 weeks ago

Selected Answer: AC

A, C is correct.

https://content.cisco.com/chapter.sjs?uri=/searchable/chapter/www.cisco.com/content/en/us/td/docs/interfaces_modules/shared_port_adapters/configuration/6500series/sipspasw/76cfvpnb.html.xml

upvoted 1 times

mitosenoriko 11 months, 2 weeks ago

A and C is correct
I checked cisco documents.

upvoted 1 times

Zizu007 11 months, 2 weeks ago

Selected Answer: AC

Correct!

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_dplane/configuration/xe-16-10/sec-ipsec-data-plane-xe-16-10-book/sec-pre-frag-vpns.html

Restrictions for Pre-Fragmentation for IPsec VPNs
Take the following information into consideration before this feature is configured:

Pre-fragmentation for IPsec VPNs operates in IPsec tunnel mode and IPsec tunnel mode with GRE, but not with IPsec transport mode.

Pre-fragmentation for IPsec VPNs configured on the decrypting router in a unidirectional traffic scenario does not improve the performance or change the behavior of either of the peers.

Pre-fragmentation for IPsec VPNs occurs before the transform is applied if compression is turned on for outgoing packets.

Pre-fragmentation for IPsec VPNs functionality depends on the egress interface crypto ipsec df-bit configuration and the incoming packet “do not fragment” (DF) bit state. See the table below.

upvoted 1 times

Question #331 *Topic 1*

Which of the following correctly describes the concept of split horizon with IP routing? (Choose two.)

A. Split horizon is a valid routing loop prevention mechanism

B. Split horizon is used to filter customer routes in an ISP network.

C. When enabled, split horizons informs the router to not advertise routes back out the same interface from where that route was originally received.

D. Split horizons cannot be disabled on WAN interfaces

E. Split horizon is not applicable to EIGRP networks

Correct Answer: *AC*

alex711 3 months, 3 weeks ago

The given answer is correct.

upvoted 1 times

mitosenoriko 11 months, 2 weeks ago

A and C is correct.

upvoted 2 times

Noproblem22 1 year ago

A and C are correct

upvoted 4 times

Question #332

Topic 1

DRAG DROP

-

Arrange the below as per the recommended steps:

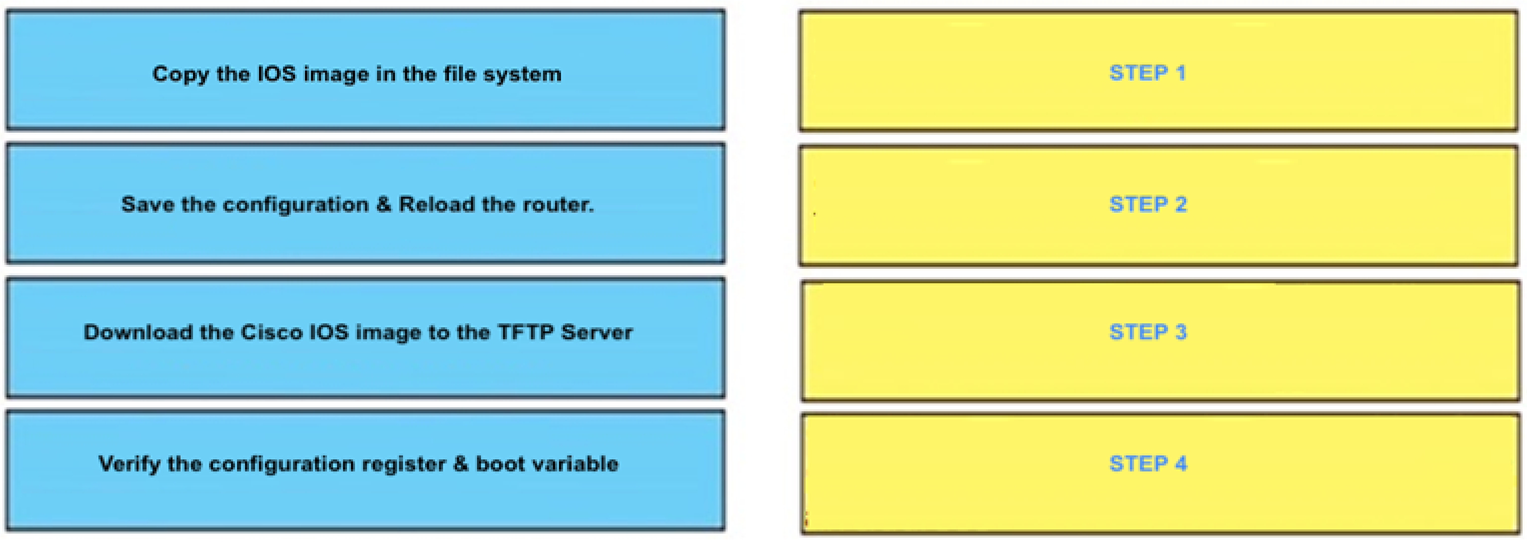

Copy the IOS image in the file system

Save the configuration & Reload the router.

Download the Cisco IOS image to the TFTP Server

Verify the configuration register & boot variable

STEP 1

STEP 2

STEP 3

STEP 4

Correct Answer:

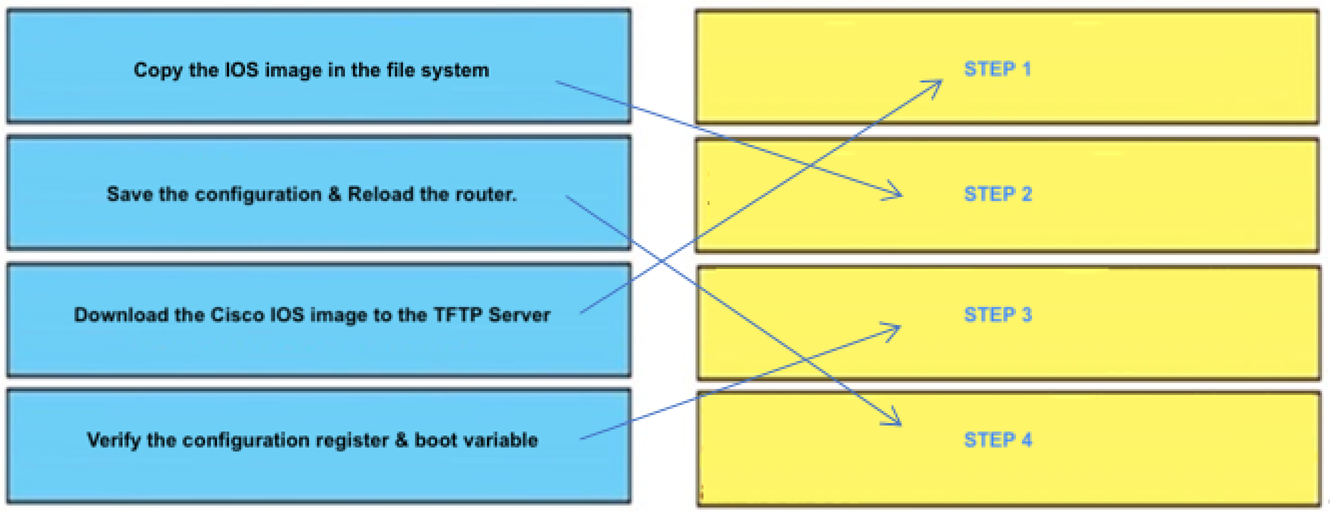

Copy the IOS image in the file system

Save the configuration & Reload the router.

Download the Cisco IOS image to the TFTP Server

Verify the configuration register & boot variable

STEP 1

STEP 2

STEP 3

STEP 4

⊟ **mitosenoriko** 11 months, 2 weeks ago

this procedure is correct.

upvoted 2 times

Question #333 *Topic 1*

A network administrator is reloading a router and during the bootup, he is getting the error message “%Error opening tftp://255.255.255.255/network-confg (Socket error)”. What command need to be applied on Cisco Router to fix this issue.

A. No service config

B. Write erase reload

C. Reload noconfirm

D. Copy run start

Correct Answer: *A*

sayed_2908 10 months, 3 weeks ago

Ans:A

https://community.cisco.com/t5/networking-knowledge-base/the-router-continually-tries-to-load-a-configuration-from-the/ta-p/3131171

upvoted 2 times

mitosenoriko 11 months, 2 weeks ago

What to do depends on the situation.
A.Probably OK.
B.Remote monitoring would be a big problem.
C.Retry this problem
D.Saving bad configurations?

upvoted 1 times

Brand 3 months, 2 weeks ago

C1-HUB(config)#no service ?
compress-config Compress the nvram configuration file
config TFTP load config files

upvoted 2 times

Question #334 Topic 1

DRAG DROP

-

The steps for configuring BGP on Cisco IOS Router:

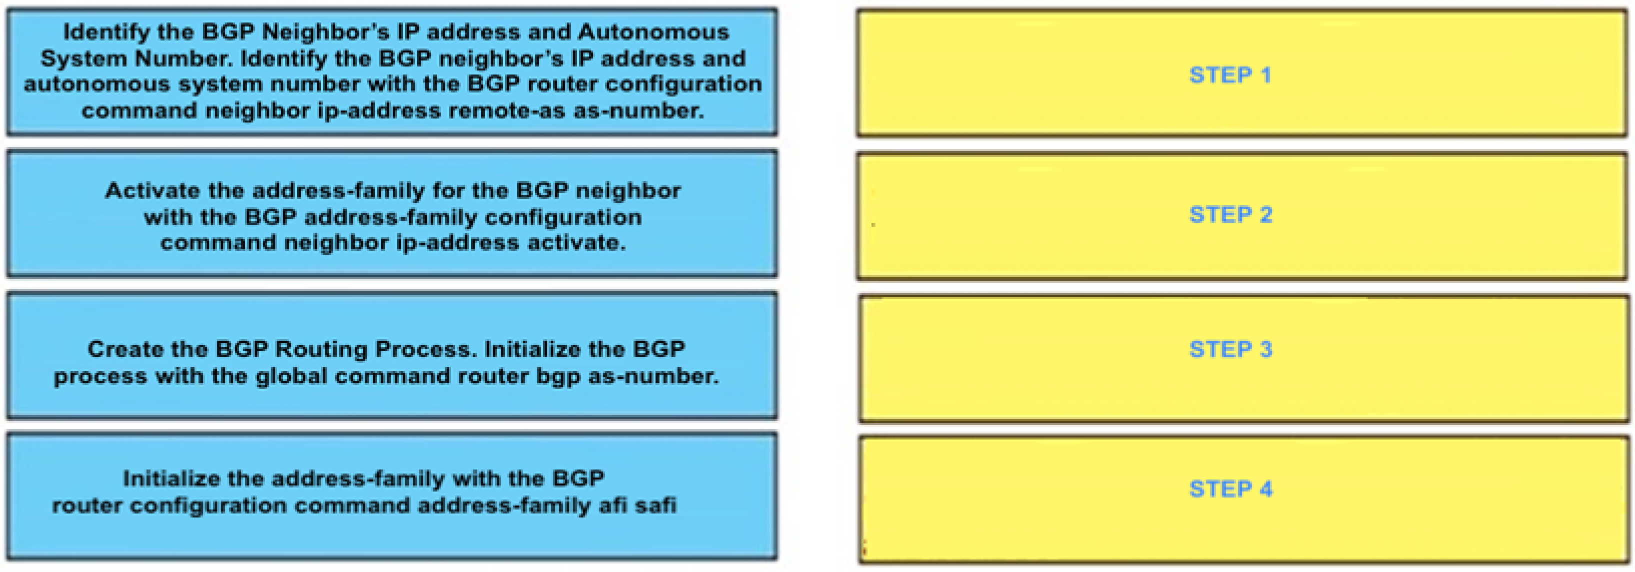

Identify the BGP Neighbor's IP address and Autonomous System Number. Identify the BGP neighbor's IP address and autonomous system number with the BGP router configuration command neighbor ip-address remote-as as-number.

Activate the address-family for the BGP neighbor with the BGP address-family configuration command neighbor ip-address activate.

Create the BGP Routing Process. Initialize the BGP process with the global command router bgp as-number.

Initialize the address-family with the BGP router configuration command address-family afi safi

Correct Answer:

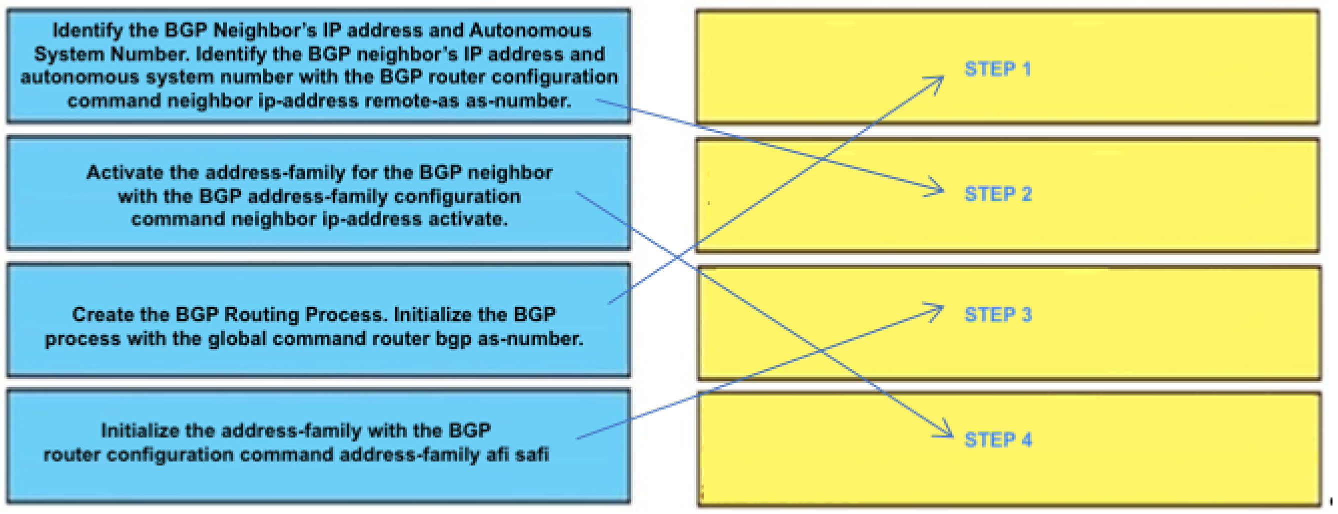

Identify the BGP Neighbor's IP address and Autonomous System Number. Identify the BGP neighbor's IP address and autonomous system number with the BGP router configuration command neighbor ip-address remote-as as-number.

Activate the address-family for the BGP neighbor with the BGP address-family configuration command neighbor ip-address activate.

Create the BGP Routing Process. Initialize the BGP process with the global command router bgp as-number.

Initialize the address-family with the BGP router configuration command address-family afi safi

inteldarvid 5 months, 1 week ago

Given Awser correct

upvoted 1 times

GReddy2323 6 months ago

IOS

The steps for configuring BGP on an IOS router are as follows:

Step 1. Create the BGP Routing Process. Initialize the BGP process with the global command router bgp as-number.

Step 2. Identify the BGP Neighbor's IP address and Autonomous System Number. Identify the BGP neighbor's IP address and autonomous system number with the BGP router configuration command neighbor ip-address remote-as as-number.

NOTE

IOS activates the IPv4 address-family by default. This can simplify the configuration in an IPv4 environment because Steps 3 and 4 are optional, but may cause confusion when working with other address families. The BGP router configuration command no bgp default ip4-unicast disables the automatic activation of the IPv4 AFI so that Steps 3 and 4 are required.

Step 3. Initialize the address-family with the BGP router configuration command address-family afi safi.

Step 4. Activate the address-family for the BGP neighbor with the BGP address-family configuration command neighbor ip-address activate.

upvoted 2 times

ellen_AA 11 months, 2 weeks ago

Given Answer is correct.
https://www.ciscopress.com/articles/article.asp?p=2756480&seqNum=5

upvoted 1 times

DUBC89x 1 year ago

Given answer is correct.
SUMMARY STEPS
enable
configure terminal
router bgp autonomous-system-number
neighbor ip-address remote-as autonomous-system-number

address-family ipv4 [unicast | multicast | vrf vrf-name]
neighbor ip-address activate
end
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-16/irg-xe-16-book/configuring-a-basic-bgp-network.html
upvoted 2 times

Question #335 Topic 1

What is the term used when it causes the packets to lose their MPLS labels including the VPN in-formation that lies in the inner MPLS Label i.e. if a packet goes through an untagged interface, the VPN information is lost and VPN sites lose connectivity?

A. Pseudowire

B. Black Hole

C. Traffic Engineering

D. Active Network Abstraction

Correct Answer: *B*

Muste 4 months ago

Selected Answer: B

probided answer is correct : One common cause of black holes in MPLS is the mismatch between the IP routing table and the label forwarding table. For example, if a router receives a packet with a label that is not in its label forwarding table, it will drop the packet without sending any error message.
upvoted 2 times

Question #336 Topic 1

An administrator wants to implement security on his company's router. Please select three options that you will use on your router to secure it. (Choose three.)

A. Control Access to the router

B. Restrict all traffic through the router

C. Restrict SNMP

D. Enable all unused services

E. Encrypt all passwords

F. Disable logging

Correct Answer: *ACE*

amadeu 11 months, 2 weeks ago

Correct: ACE
upvoted 3 times

Question #337 *Topic 1*

An administrator is setting up a DMVPN tunnel between their offices and he is getting below output when he is running the command “show crypto isakmp sa”:

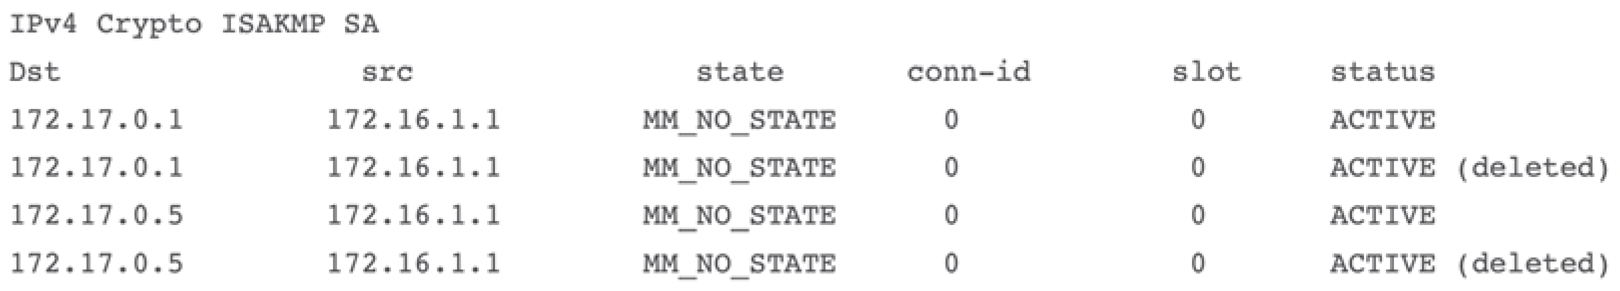

```
IPv4 Crypto ISAKMP SA
Dst            src            state        conn-id      slot     status
172.17.0.1     172.16.1.1     MM_NO_STATE    0            0      ACTIVE
172.17.0.1     172.16.1.1     MM_NO_STATE    0            0      ACTIVE (deleted)
172.17.0.5     172.16.1.1     MM_NO_STATE    0            0      ACTIVE
172.17.0.5     172.16.1.1     MM_NO_STATE    0            0      ACTIVE (deleted)
```

What command will you run to identify the issue?

A. Debug ip icmp

B. Debug crypto isakmp

C. Debug crypto ipsec sa

D. Debug ssh

Correct Answer: *B*

HarwinderSekhon 4 months, 1 week ago

Selected Answer: B

If there's an issue with the IKE phase of the tunnel establishment, the appropriate command to run for further troubleshooting would be:

B. Debug crypto isakmp

This command enables detailed debugging of the IKE protocol, which would help the administrator to identify the issues related to the IKE phase of the VPN tunnel establishment.

upvoted 2 times

Question #338 *Topic 1*

A company is looking to implement VPN between their Head Quarter and over 100+ Branch Offices. They are looking for a solution that:

1. Reduces deployment complexity
2. Simplifies branch communications
3. Offers branch to branch connectivity.
4. Is cost effective
5. Offers strong encryption

Select the best option from the below options that you would recommend to implement.

A. MPLS

B. IPSEC

C. DMVPN

D. GRE

Correct Answer: *C*

Question #339 *Topic 1*

You have a DNA center deployed in your environment. Which feature of the DNA Center will you use for system-guided as well as self-guided troubleshooting.

A. Assurance

B. Automation

C. Zero Trust

D. Discovery

Correct Answer: *A*

Malasxd 7 months, 2 weeks ago

Selected Answer: A

Agree. Answer "A"

upvoted 1 times

GReddy2323 9 months, 2 weeks ago

Selected Answer: A

• Provides both system-guided as well as self-guided troubleshooting. For a large number of issues, Assurance provides a system-guided approach, where multiple Key Performance Indicators (KPIs) are correlated, and the results from tests and sensors are used to determine the root cause of a problem, after which possible actions are provided to resolve the problem. The focus is on highlighting the issue rather than monitoring data. Quite frequently, Assurance performs the work of a Level 3 support engineer.

https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center-assurance/1-3-3-0/b_cisco_dna_assurance_1_3_3_0_ug/b_cisco_dna_assurance_1_3_2_0_chapter_01.pdf

upvoted 1 times

Question #340 *Topic 1*

DRAG DROP

-

You are logged in to the DNA Center Client Health Dashboard. Under the client health, you see some color-coded fields that reflects the health status of the client devices. Drag the health scores on the left to their respective colors in the right.

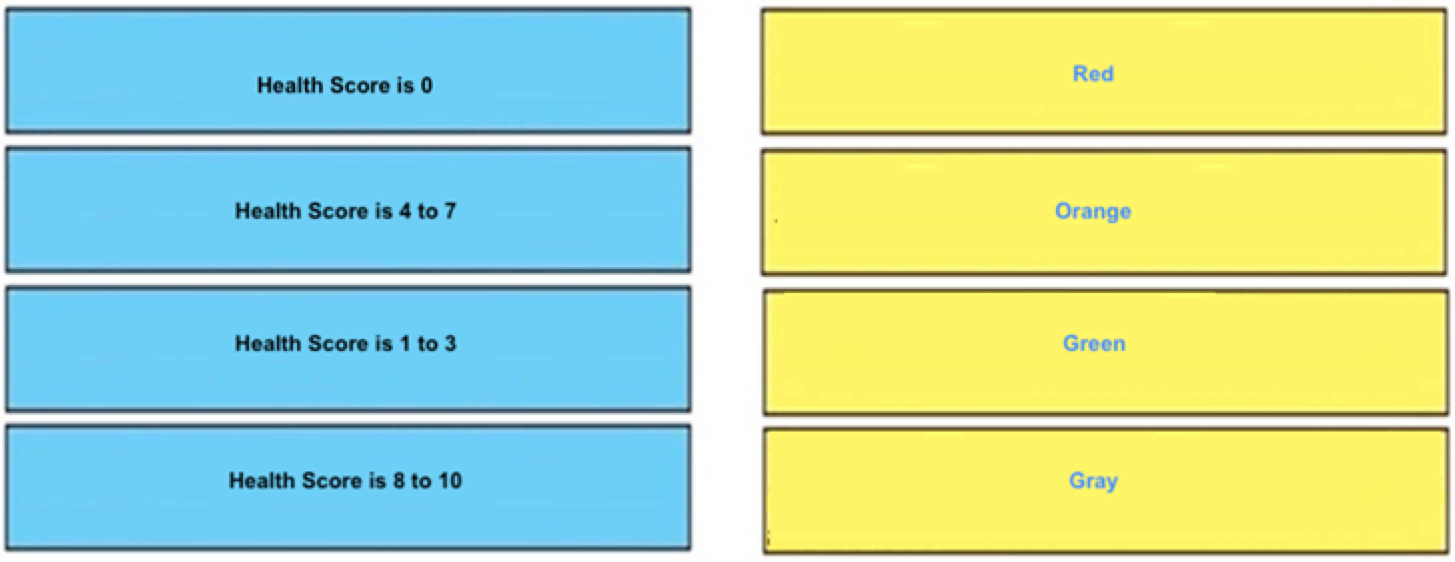

Health Score is 0

Health Score is 4 to 7

Health Score is 1 to 3

Health Score is 8 to 10

Red

Orange

Green

Gray

Correct Answer:

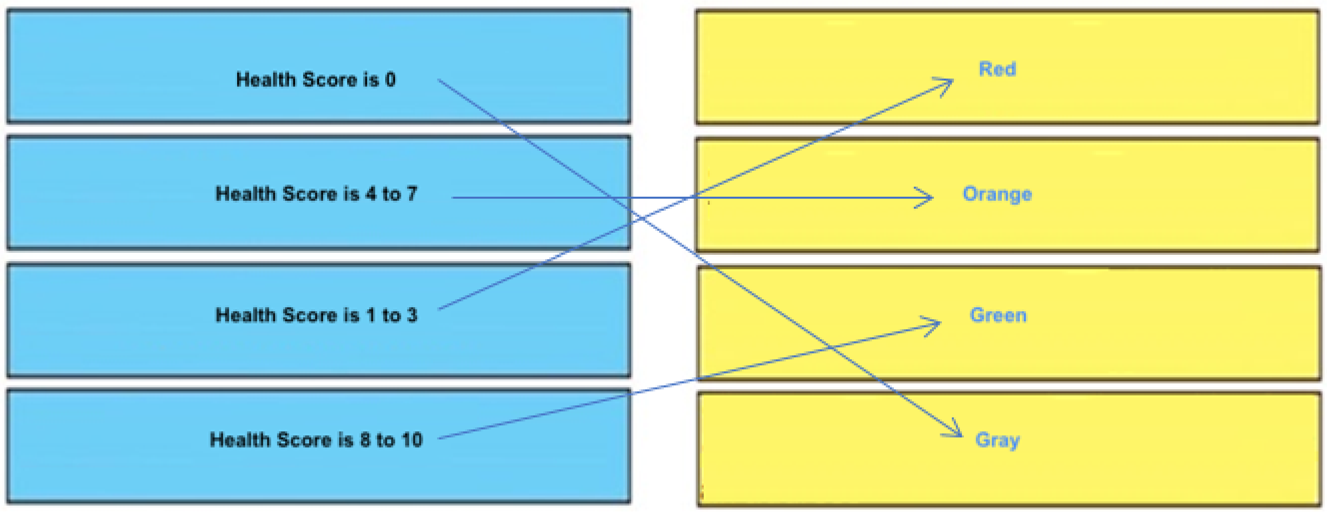

Health Score is 0

Health Score is 4 to 7

Health Score is 1 to 3

Health Score is 8 to 10

Red

Orange

Green

Gray

DUBC89x 1 year ago

Given answer is correct.
The color of the health score represents its severity. The health is measured on a scale of 1 to 10, with 10 being the best score, and a score of 0 indicating that the client is inactive.
https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center-assurance/1-3-1-0/b_cisco_dna_assurance_1_3_1_0_ug/b_cisco_dna_assurance_1_3_1_0_chapter_0101.html

upvoted 2 times

Question #341 *Topic 1*

Out of the below options regarding DMVPN & FLEXVPN, select the correct one.

A. FlexVPN uses a new key management protocol – IKEv2, while most traditional DMVPN networks use IKEv1

B. FlexVPN uses a new key management protocol – IKEv1, while most traditional DMVPN networks use IKEv2

C. With FlexVPN there's multiple standard way of NHRP and routing protocols operations as opposed to 1 phase of DMVPN

D. Flex VPN & DMVPN both are supported only on Firewalls.

Correct Answer: *A*

⊟ **inteldarvid** 5 months, 1 week ago

https://community.cisco.com/t5/network-security/what-is-the-difference-between-dmvpn-and-flexvpn/td-p/3438913

upvoted 1 times

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: A

yes, A correct

upvoted 1 times

⊟ **mitosenoriko** 11 months, 2 weeks ago

A is correct

upvoted 3 times

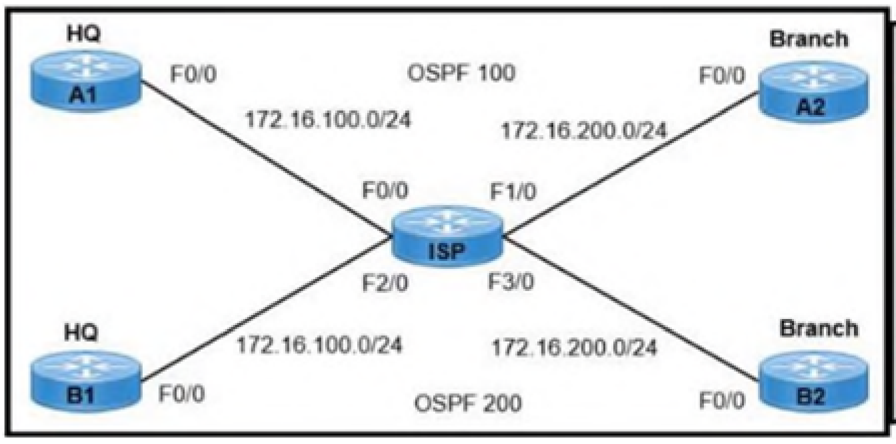

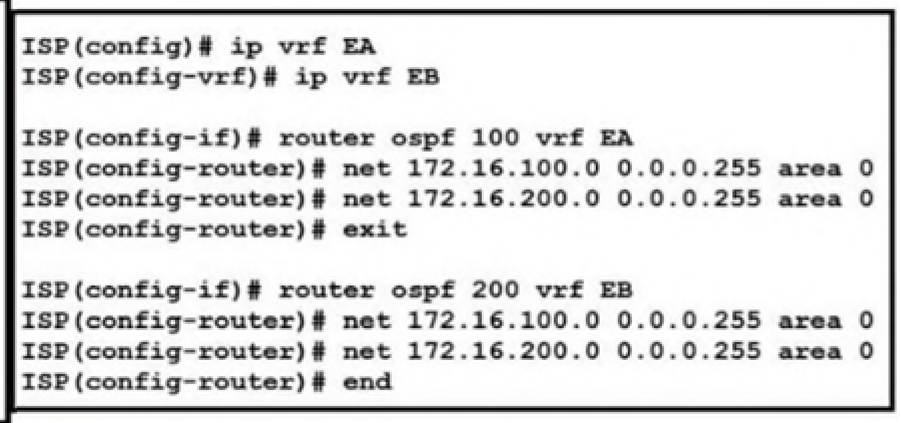

```
ISP(config)# ip vrf EA
ISP(config-vrf)# ip vrf EB

ISP(config-if)# router ospf 100 vrf EA
ISP(config-router)# net 172.16.100.0 0.0.0.255 area 0
ISP(config-router)# net 172.16.200.0 0.0.0.255 area 0
ISP(config-router)# exit

ISP(config-if)# router ospf 200 vrf EB
ISP(config-router)# net 172.16.100.0 0.0.0.255 area 0
ISP(config-router)# net 172.16.200.0 0.0.0.255 area 0
ISP(config-router)# end
```

Refer to the exhibit. A network engineer is provisioning end-to-end traffic service for two different enterprise networks with these requirements:

• The OSPF process must differ between customers on HQ and Branch office routers, and adjacencies should come up instantly.

• The enterprise networks are connected with overtapping networks between HQ and a Branch office.

Which configuration meets the requirements for a customer site?

A. ISP(config-if)#int f1/0 -
ISP(config-if)#ip vrf forwarding EA
ISP(config-if)#description TO->EA2_Branch
ISP(config-if)#ip add 172.16.200.2 255.255.255.0

ISP(config-if)#no shut -

B. ISP(config-vrf)#int f0/0 -
ISP(config-if)#ip vrf forwarding EB
ISP(config-if)#description TO->EB1_Branch
ISP(config-if)#ip add 172.16.100.2 255.255.255.0

ISP(config-if)#no shut -

C. ISP(config)#int f2/0 -
ISP(config-if)#ip vrf forwarding EA
ISP(config-if)#description TO->EA1_HQ
ISP(config-if)#ip address 172.16.100.2 255.255.255.0

ISP(config-if)#no shut -

D. ISP(config-if)#int f3/0 -
ISP(config-if)#ip vrf forwarding EA
ISP(config-if)#description TO->EA2_Branch
ISP(config-if)#ip address 172.16.200.2 255.255.255.0
ISP(config-if)#no shut

Correct Answer: *A*

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: A

A correct

upvoted 1 times

⊟ **Almylle** 5 months, 4 weeks ago

Selected Answer: A

A is correct, the other alternatives all have the worng ip vrf forwarding for the interfaces.

upvoted 1 times

⊟ **6dd4aa0** 8 months, 2 weeks ago

A is correct.

Took me a long while to figure out.

The hint is at the OSPF commands. Followed by, the figure where it shows the two OSPF regions must tally with VRF. VRF EA points to OSPF 100 and VRF EB points to OSPF 200.

Int FA0/0 and Int FA1/0 points to VRF EA, based on the OSPF 100.
Int FA2/0 and Int FA3/0 points to VRF EB, based on the OSPF 200.
upvoted 2 times

⊟ **ellen_AA** 11 months, 2 weeks ago
All ISP interfaces point to the wrong VRF, except A
upvoted 1 times

⊟ **DUBC89x** 1 year ago
I think this is not the complete picture.
upvoted 1 times

⊟ **[Removed]** 1 year ago
The answer appears to be correct. Look at the rest of the interface assignments compared to the VRF they are in. A is the only one that matches interface to the proper VRF.
upvoted 2 times

Question #343 *Topic 1*

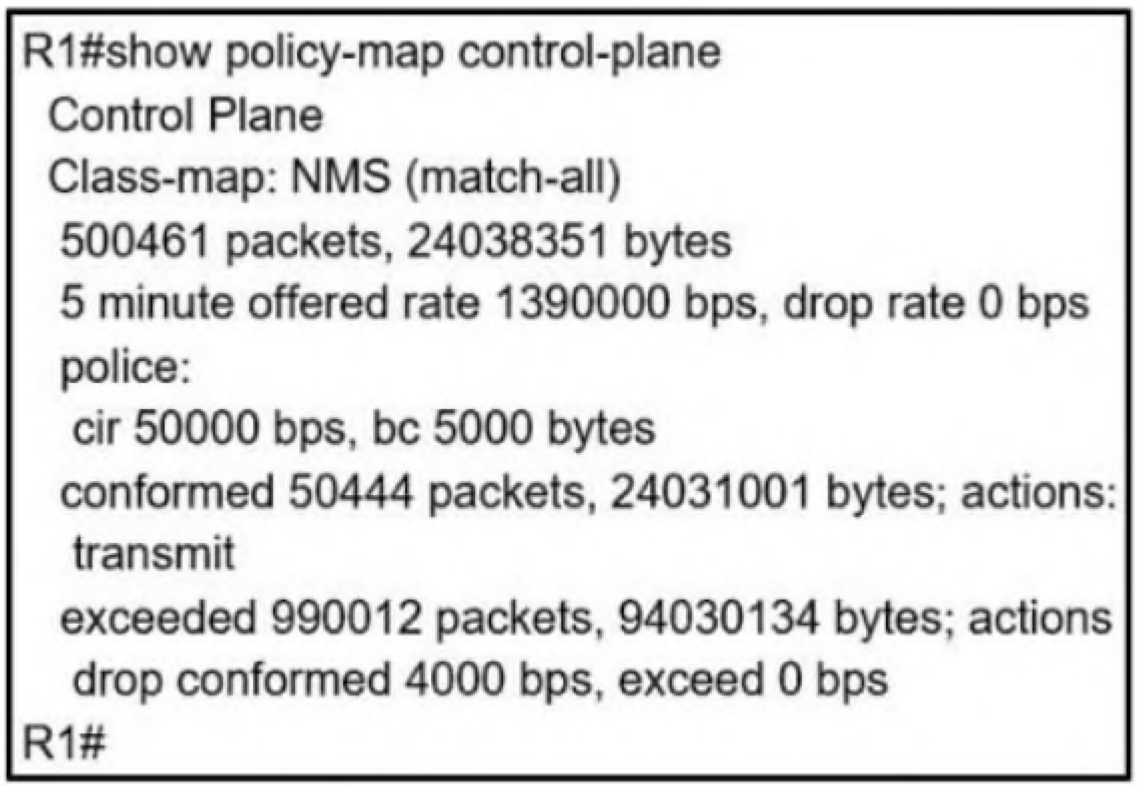

```
R1#show policy-map control-plane
  Control Plane
  Class-map: NMS (match-all)
   500461 packets, 24038351 bytes
   5 minute offered rate 1390000 bps, drop rate 0 bps
   police:
    cir 50000 bps, bc 5000 bytes
   conformed 50444 packets, 24031001 bytes; actions:
    transmit
   exceeded 990012 packets, 94030134 bytes; actions
    drop conformed 4000 bps, exceed 0 bps
R1#
```

Refer to the exhibit. A company is evaluating multiple network management system tools. Trending graphs generated by SNMP data are returned by the NMS and appear to have multiple gaps. While troubleshooting the issue, an engineer noticed the relevant output. Which action resolves the gaps in the graphs?

A. Remove the class map NMS from being part of control plane policing.

B. Configure the CIR rate to a lower value that accommodates all the NMS tools.

C. Remove the exceed-rate command in the class map.

D. Separate the NMS class map in multiple class maps based on the specific protocols with appropriate CoPP actions.

Correct Answer: *D*

⊟ **[Removed]** 4 months, 3 weeks ago
Selected Answer: A
Shouldn't it be applied to management plane for NMS applications?
upvoted 1 times

⊟ **guy276465281819372** 5 months ago
both A and D would work, Don't see how this is an appropriate question...
upvoted 1 times

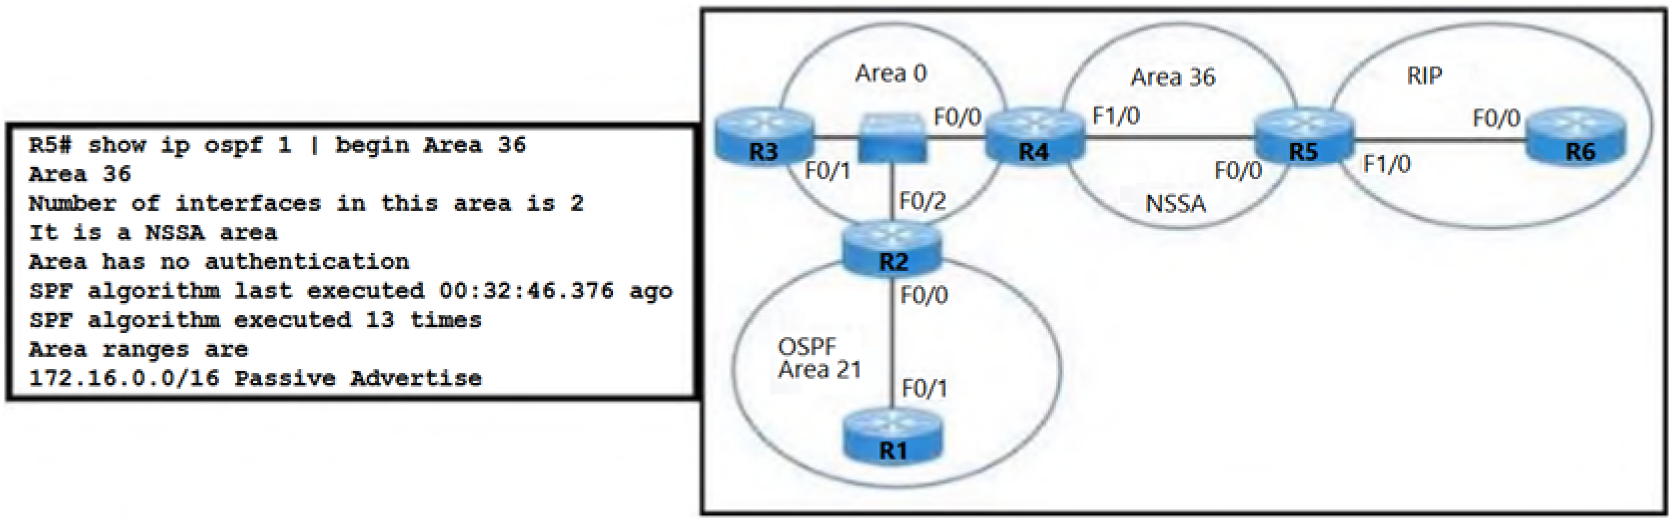

```
R5# show ip ospf 1 | begin Area 36
Area 36
Number of interfaces in this area is 2
It is a NSSA area
Area has no authentication
SPF algorithm last executed 00:32:46.376 ago
SPF algorithm executed 13 times
Area ranges are
172.16.0.0/16 Passive Advertise
```

Refer to the exhibit. The network engineer configured the summarization of the RIP routes into the OSPF domain on R5 but still sees four different 172.16.0.0/24 networks on R4. Which action resolves the issue?

A. R5(config)#router ospf 99 -
R5(config-router)#network 172.16.0.0 0.255.255.255 area 56
R5(config-router)#area 56 range 172.16.0.0 255.255.255.0

B. R5(config)#router ospf 1 -

R5(config-router)#no area -
R5(config-router)#summary-address 172.16.0.0 255.255.252.0

C. R4(config)#router ospf 1 -

R4(config-router)#no area -
R4(config-router)#summary-address 172.16.0.0 255.255.252.0

D. R4(config)#router ospf 99 -
R4(config-router)#network 172.16.0.0 0.255.255.255 area 56
R4(config-router)#area 56 range 172.16.0.0 255.255.255.0

Correct Answer: *B*

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: B

100% Correct in ASBR:(summary address)
http://ithitman.blogspot.com/2015/04/configuring-ospf-area-range-vs-summary.html

upvoted 1 times

⊟ **Almylle** 5 months, 4 weeks ago

Selected Answer: B

The summary address only can be configurated in the ASBR routers in OSPF, in this case is only the R5, and u can discard the other alternative because they are trying other ospf process, and the ospf process has to be 1

upvoted 1 times

⊟ **sayed_2908** 10 months, 3 weeks ago

Selected Answer: B

R4 is an ABR. R5 is an ASBR.
Area range is used on ABR and summary addres on ASBR.

Ans is B.

upvoted 1 times

⊟ **GReddy2323** 10 months, 3 weeks ago

Can someone please explain?

upvoted 1 times

⊟ **MonzaInA** 7 months, 4 weeks ago

Exclude area range because is ASBR and not ABR
Than the router is R5 so B

upvoted 1 times

Question #345

Topic 1

What is the minimum time gap required by the local system before putting a BFD control packet on the wire?

A. Desired Min TX Interval

B. Detect Mult

C. Required Min RX Interval

D. Required Min Echo RX Interval

Correct Answer: *A*

ZamanR 1 week, 4 days ago

A is Answer
Desired Min TX Interval: This is the minimum interval, in microseconds, that the local system would like to use when transmitting BFD Control packets, less any jitterapplied. The value zero is reserved.

Required Min Echo RX Interval: This is the minimum interval, in microseconds, between received BFD Echo packets that this system is capable of supporting, less anyjitter applied by the sender. If this value is zero, the transmitting system does not support the receipt of BFD Echo packets.

Reference: https://tools.ietf.org/html/rfc5880

upvoted 1 times

HarwinderSekhon 4 months, 1 week ago

Selected Answer: A

This value specifies the minimum interval, in microseconds, that the local system would like to use when transmitting BFD control packets. It helps in controlling how frequently the BFD control packets are sent, allowing for the tuning of BFD's detection time.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: A

yes. option A:
https://www.cisco.com/en/US/technologies/tk648/tk365/tk480/technologies_white_paper0900aecd80244005.html

upvoted 1 times

HamzaBadar 5 months, 2 weeks ago

Selected Answer: A

Answer is A.

upvoted 1 times

dax_bux 7 months, 1 week ago

Same as Question #42

upvoted 1 times

Lorenzzz987 5 months, 1 week ago

Not quite, the minimal interval that the local system like to use is the Desired Minimum TX interval. Question 42 expects the minimum interval between received BFD control packets that this system is capable of supporting, hence Min RX Interval.

upvoted 2 times

pitcholo 10 months, 3 weeks ago

Desired Min TX Interval : This is the minimum interval, in microseconds, that the local system would like to use when transmitting BFD Control packets.

Required Min RX Interval : This is the minimum interval, in microseconds, between received BFD Control packets that this system is capable of supporting.

Answer is A.

upvoted 3 times

Supawit_t 1 year ago

Selected Answer: A

Select A
https://www.ietf.org/rfc/rfc5880.txt
Desired Min TX Interval

This is the minimum interval, in microseconds, that the local
system would like to use when transmitting BFD Control packets,
less any jitter applied (see section 6.8.2). The value zero is
reserved.

upvoted 1 times

Netking 1 year ago

Selected Answer: A

The given answer is correct

upvoted 1 times

msama 1 year ago

Selected Answer: C

Same as Question #42

upvoted 2 times

DUBC89x 1 year ago

Selected Answer: C

Required Min RX Interval: This is the minimum interval, in microseconds, between received BFD Control packets that this system is capable of supporting, less any jitter applied by the sender

Required Min RX Interval: This is the minimum interval, in microseconds, between received BFD Control packets that this system is capable of supporting, less any jitter applied by the sender

"https://www.networkworld.com/article/2222661/bidirectional-forwarding-detection--bfd----a-little-about-timers.html#:~:text=Desired%20Min%20TX%20Interval%3A%20This,The%20value%20zero%20is%20reserved."

The question asked "required" not desired.

upvoted 2 times

Question #346 *Topic 1*

What must be configured by the network engineer to circumvent AS_PATH loop prevention mechanism in IP/VPN Hub and Spoke deployment scenarios?

A. Use allowas-in at the PE_Hub.

B. Use allowas-in and as-override at all PEs.

C. Use allowas-in and as-override at the PE_Hub.

D. Use as-override at the PE_Hub.

Correct Answer: *D*

ZamanR 5 days, 2 hours ago

D is correct

upvoted 1 times

HarwinderSekhon 4 months, 1 week ago

Selected Answer: D

AS-Override:

Purpose: This is used to replace occurrences of the AS number of the provider edge (PE) router with the AS number of the remote AS, typically in the context of BGP VPNs.
Typical Use: This is often used by service providers in a BGP/MPLS VPN to prevent loops in the customer's network when the customer is using the same AS number at multiple sites.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: D

yes, correct option "D"

upvoted 2 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: D

AS override: this can be configured on the PE routers, the AS number will be replaced with the AS number from the service provider.
Allow-AS in: this can be configured on the CE routers which tells them to accept prefixes with their own AS number in the AS path.
https://networklessons.com/bgp/mpls-layer-3-vpn-bgp-override

https://content.cisco.com/chapter.sjs?uri=/searchable/chapter/content/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3s/irg-xe-3s-book/bgp-as-override-split-horizon.html.xml

upvoted 3 times

6dd4aa0 8 months, 2 weeks ago

Selected Answer: D

While former is used by PE to modify the AS Number in AS Path so that prefix is not dropped, the latter is implemented in CE device to introduce an exception in BGP AS path loop prevention mechanism.

Hub should be seen as a provider edge, and Spoke as a customer edge.

"If the customer requires to keep minimal configuration at CE side and let the provider perform the BGP routing control, the best approach will be to use “As-Override”."

Hence, the Hub is provider edge, it should be configured with AS-Override

upvoted 3 times

Titini 10 months ago

Selected Answer: C

We need both commands to prevent loops

upvoted 2 times

DUBC89x 1 year ago

Well after reading more I believe D is correct. Since the answers are only on the provider end.

upvoted 4 times

DUBC89x 1 year ago

Selected Answer: A

1. AS Override :
Its feature allows a provider edge (PE) router to change private autonomous system used by customer edge (CE) device on an external BGP session running on a VPN routing and forwarding access link. The private AS number is changed to PEAS number.

2. Allowas In :
This feature allows for routes to be received and processed even if router detects its own ASN in AS-Path. A router discards BGP network prefixes if it sees its ASN in AS-Path as a loop prevention mechanism.
"https://www.geeksforgeeks.org/difference-between-as-override-and-allowas-in/#:~:text=2.,as%20a%20loop%20prevention%20mechanism."

upvoted 1 times

Question #347 Topic 1

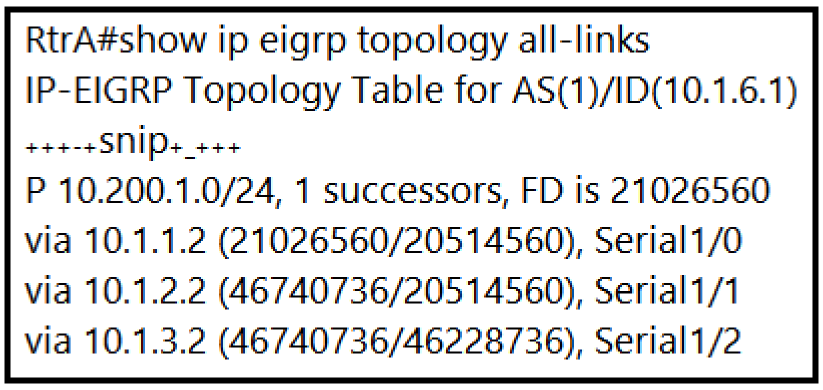

```
RtrA#show ip eigrp topology all-links
IP-EIGRP Topology Table for AS(1)/ID(10.1.6.1)
++++snip+_+++
P 10.200.1.0/24, 1 successors, FD is 21026560
via 10.1.1.2 (21026560/20514560), Serial1/0
via 10.1.2.2 (46740736/20514560), Serial1/1
via 10.1.3.2 (46740736/46228736), Serial1/2
```

Refer to the exhibit. Which action makes 10.1.3.2 the feasible successor to reach 10.200.1.0/24 for location S42T431E64F51?

A. Increase path bandwidth higher than 10.1.1.2 and lower than 10.1.2.2 between RtrA and the destination.

B. Increase path bandwidth lower than 10.1.1.2 and lower than 10.1.2.2 between RtrA and the destination.

C. Increase path bandwidth higher than 10.1.2.2 and lower than 10.1.1.2 between RtrA and the destination.

D. Increase path bandwidth higher than 10.1.2.2 and higher than 10.1.1.2 between RtrA and the destination.

Correct Answer: *D*

ccnptoppler34 Highly Voted 11 months, 2 weeks ago

Selected Answer: C

I don't believe D is correct the question asks to make the the route to 10.1.3.2 the feasible successor not the successor

upvoted 12 times

inteldarvid Most Recent 5 months, 1 week ago

Selected Answer: C

yesy optio c correct

upvoted 1 times

pyrokar 6 months, 4 weeks ago

Selected Answer: C

10.1.3.2 can't be the feasible successor since feasibility condition is not met. Apart from that, it should be C.

upvoted 2 times

HungarianDish 7 months, 1 week ago

Selected Answer: C

Correct is "C".
https://networklessons.com/eigrp/eigrp-wide-metrics

Further info:
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html
If it is necessary to influence the path EIGRP chooses, always use delay to do so.

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/13673-14.html
There are two reasons you must not to use the bandwidth to influence EIGRP paths:
-If you change the bandwidth it can have impact beyond the effect to the EIGRP metrics.
-When you lower the bandwidth it can prevent EIGRP neighbors that get hello packets because it throttles back.

If you change the delay it does not impact other protocols nor does it cause EIGRP to throttle back.

upvoted 3 times

forccnp 9 months, 1 week ago

Selected Answer: C

D will make 10.1.3.2 the successor

upvoted 2 times

Titini 10 months, 1 week ago

Selected Answer: C

Answer is C

upvoted 1 times

sayed_2908 10 months, 3 weeks ago

Selected Answer: C

D will make 10.1.3.2 the successor and we don't want that.

upvoted 1 times

toto89 11 months, 2 weeks ago

Selected Answer: C

I also think C is the right answer, minimum path bandwidth lower than 10.1.1.2 and higher than 10.1.2.2 should do it.
But again a weird question from Cisco, we don't know anything about the cumulative delay and other parameters such as the number of hops between the local router and the destination for each path.
B could also do it since all we want here is meeting the feasibility condition for the 10.1.3.2 path. But C is the more likely to be true.

upvoted 1 times

Question #348 *Topic 1*

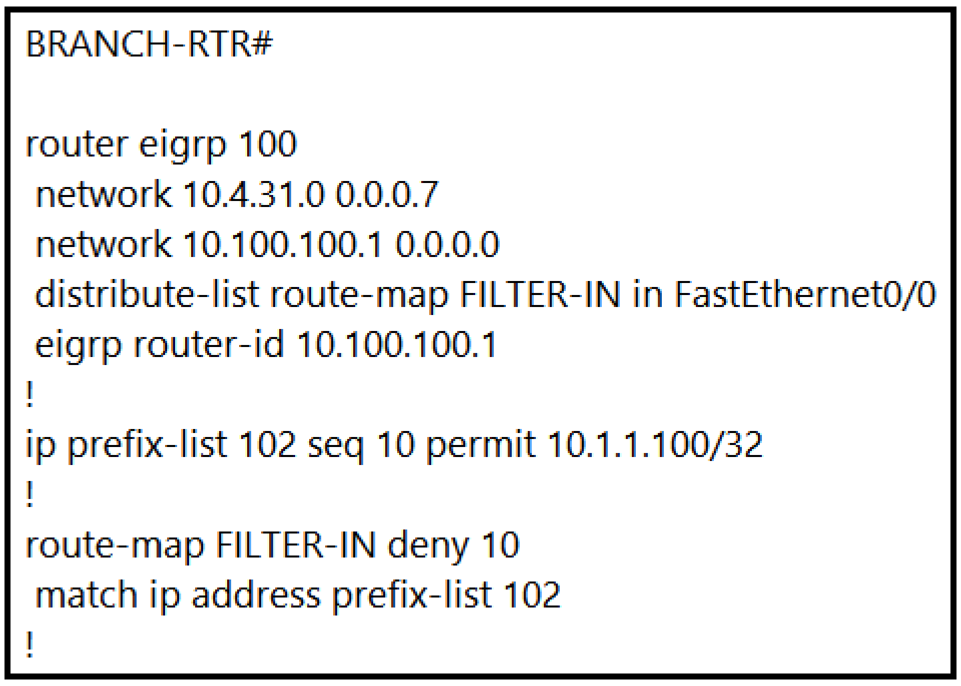

```
BRANCH-RTR#

router eigrp 100
 network 10.4.31.0 0.0.0.7
 network 10.100.100.1 0.0.0.0
 distribute-list route-map FILTER-IN in FastEthernet0/0
 eigrp router-id 10.100.100.1
!
ip prefix-list 102 seq 10 permit 10.1.1.100/32
!
route-map FILTER-IN deny 10
 match ip address prefix-list 102
!
```

Refer to the exhibit A junior engineer updated a branch router configuration. Immediately after the change, the engineer receives calls from the help desk that branch personnel cannot reach any network destinations. Which configuration restores service and continues to block 10.1.1.100/32?

A. route-map FILTER-IN deny 5

B. ip prefix-list 102 seq 15 permit 0.0.0.0/32 le 32

C. route-map FILTER-IN permit 20

D. ip prefix-list 102 seq 5 permit 0.0.0.0/32 le 32

Correct Answer: *C*

[Removed] 4 months, 3 weeks ago

Selected Answer: C

I got tricked by not reading properly.

Remember that a prefix list and a route map have an implicit deny at the end of their entries.
The prefix list is permitting the host prefix so that it can be processed by the Route map. And it purposely leaves the implicit deny because the Route map only processes that host to deny it. However, the Route map implicitly denies everything else preventing other routes from being installed in the RIB

upvoted 2 times

forccnp 10 months ago

Selected Answer: C

Given answer is correct

upvoted 2 times

Question #349 *Topic 1*

What does the MP-BGP OPEN message contain?

A. the version number and the AS number to which the router belongs

B. IP routing information and the AS number to which the router belongs

C. NLRI, path attributes, and IP addresses of the sending and receiving routers

D. MPLS labels and the IP address of the router that receives the message

Correct Answer: *A*

⊟ **DUBC89x** Highly Voted 1 year ago

Selected Answer: A

The OPEN message contains the BGP version number, ASN of the originating router, Hold Time, BGP Identifier, and other optional parameters that establish the session capabilities.

upvoted 6 times

Question #350 Topic 1

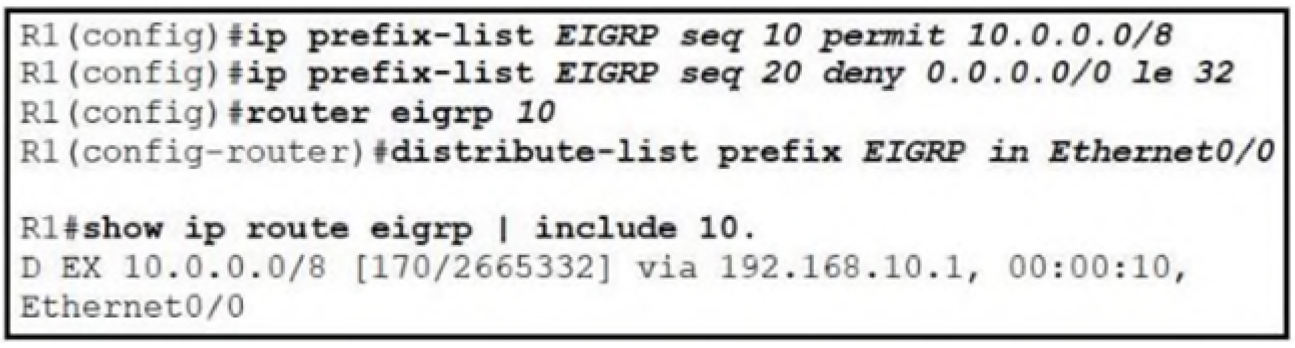

```
R1(config)#ip prefix-list EIGRP seq 10 permit 10.0.0.0/8
R1(config)#ip prefix-list EIGRP seq 20 deny 0.0.0.0/0 le 32
R1(config)#router eigrp 10
R1(config-router)#distribute-list prefix EIGRP in Ethernet0/0

R1#show ip route eigrp | include 10.
D EX 10.0.0.0/8 [170/2665332] via 192.168.10.1, 00:00:10,
Ethernet0/0
```

Refer to the exhibit. An engineer applies a prefix-list filter that filters most of the network 10 prefixes instead of allowing them. Which action resolves the issue?

A. Modify the ip prefix-list EIGRP seq 20 permit 10.0.0.0/8 ge 9 command.

B. Modify the ip prefix-list EIGRP seq 10 permit 10.0.0.0/8 le 9 command.

C. Modify the ip prefix-list EIGRP seq 20 permit 0.0.0.0/0 le 32 command.

D. Modify the ip prefix-list EIGRP seq 10 permit 10.0.0.0/8 le 32 command.

Correct Answer: *D*

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: D

yes option correct "D"

upvoted 1 times

⊟ **ClaudeYun** 8 months, 2 weeks ago

Selected Answer: D

I would say the question describes the issue is engineer intend to permit 10.0.0.0/8 networks but only 10.0.0.0/8 left, other precise subnet like 10.1.1.0/24 missed. how to fix it.
then we can say option D is the correct answer.

upvoted 4 times

⊟ **[Removed]** 4 months, 3 weeks ago

Yeah. With Cisco you have to smoke the same shit the exam writer did to understand their fucked up wording. Like the outsourced the job to a non native English speaker

upvoted 3 times

⊟ **Lilienen** 10 months, 2 weeks ago

Selected Answer: D

Correct answer: D
ip prefix-list EIGRP seq 10 permit 10.0.0.0/8 le 32

@ellen_AA:
I think the goal is to allow the prefixes contained in 10.0.0.0/8, not deny them.
So we need to add 'le 32' statement within the permit sequence.

upvoted 1 times

⊟ **GReddy2323** 9 months, 2 weeks ago

I am confused, doesn't the question say the engineer is filtering most of the 10 networks instead of allowing them? I am very confused with this question.

upvoted 2 times

⊟ **ellen_AA** 11 months, 2 weeks ago

None is correct!
We should have deny statement instead!

#ip prefix-list seq 10 deny 10.0.0.0/8 ge 9 (denies 10 prefixes)
#ip prefix-list seq 10 permit 0.0.0.0/0 le 32 (allows eveything else)

upvoted 2 times

Question #351 *Topic 1*

How is a preshared key “Test” for all the remote VPN routers configured in a DMVPN using GRE over IPsec set up?

A. authentication pre-share Test address 0.0.0.0 0.0.0.0

B. set pre-share Test address 0.0.0.0 0.0.0.0

C. crypto ipsec key Test address 0.0.0.0 0.0.0.0

D. crypto isakmp key Test address 0.0.0.0 0.0.0.0

Correct Answer: *D*

⊟ **HarwinderSekhon** 4 months, 1 week ago

Selected Answer: D

```
vIOS(config)#crypto isakmp ?
aggressive-mode Disable ISAKMP aggressive mode
client Set client configuration policy
default ISAKMP default policy
disconnect-revoked-peers Disconnect Crypto Session with Revoked Peer
enable Enable ISAKMP
fragmentation IKE Fragmentation enabled if required
identity Set the identity which ISAKMP will use
invalid-spi-recovery Initiate IKE and send Invalid SPI Notify
keepalive Set a keepalive interval for use with IOS peers
key Set pre-shared key for remote peer
```

upvoted 2 times

⊟ **forccnp** 10 months ago

Selected Answer: D

Given answer is correct

upvoted 2 times

⊟ **Titini** 10 months, 1 week ago

Selected Answer: D

D is correct

upvoted 2 times

Question #352 *Topic 1*

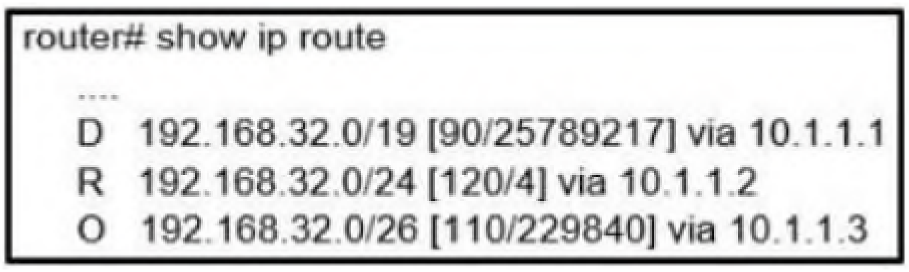

```
router# show ip route
  ....
  D   192.168.32.0/19 [90/25789217] via 10.1.1.1
  R   192.168.32.0/24 [120/4] via 10.1.1.2
  O   192.168.32.0/26 [110/229840] via 10.1.1.3
```

Refer to the exhibit. An engineer is trying to get 192.168.32.100 forwarded through 10.1.1.1, but it was forwarded through 10.1.1.2. What action forwards the packets through 10.1.1.1?

A. Configure EIGRP to receive 192.168.32.0 route with lower metric.

B. Configure EIGRP to receive 192.168.32.0 route with lower admin distance.

C. Configure EIGRP to receive 192.168.32.0 route with longer prefix than /19.

D. Configure EIGRP to receive 192.168.32.0 route with equal or longer prefix than /24.

Correct Answer: *D*

HungarianDish 6 months, 3 weeks ago

Selected Answer: D

ip prefix-list IntoEIGRP seq 10 permit 192.168.32.0/19 le /24
-> /24, 23, 22, 21, 20, 19

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

Please ignore previous post, got mixed up.
ip prefix-list IntoEIGRP seq 10 permit 192.168.32.0/19 ge /24
-> /24, 25, 26...

upvoted 1 times

MEDO95 7 months, 1 week ago

Nostalgic question from CCNA days

upvoted 3 times

drxz 7 months, 3 weeks ago

Selected Answer: D

Answer is correct. Route takes the more specific route which is a /24, so it goes trough 1.1.1.2 (RIP)
If you want it to route trough eigrp (with AD 90) you need to get the same (/24) or bigger. then it'll look for AD instead.

upvoted 1 times

juliop 9 months, 2 weeks ago

Hello,can you share me the PDF with all questions, The website is down forever(Support tell me)

upvoted 3 times

imigr 9 months, 2 weeks ago

Do we know what happend? why is it down? I will be more than happy if someone can share the pdf file.

upvoted 1 times

forccnp 9 months, 1 week ago

IT's up now(So informative)

upvoted 1 times

GReddy2323 9 months, 2 weeks ago

I don't understand this question. Could someone please explain?

upvoted 1 times

Question #353 *Topic 1*

What is a characteristic of IPv6 RA Guard?

A. It filters rogue RA broadcasts from connected hosts.

B. It is supported on the egress direction of the switch.

C. RA messages are allowed from the host port to the switch.

D. It is unable to protect tunneled traffic.

Correct Answer: *A*

Zizu007 Highly Voted 👍 11 months, 2 weeks ago

Selected Answer: D

there is no "Broadcast" in IPv6.

upvoted 9 times

inteldarvid Most Recent 5 months, 1 week ago

Selected Answer: D

option correct is "D"

upvoted 1 times

keesu 6 months ago

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-16/ip6f-xe-16-book/ip6-ra-guard.html
says
//The IPv6 RA Guard feature does not offer protection in environments where IPv6 traffic is tunneled.//

upvoted 1 times

forccnp 9 months, 1 week ago

Selected Answer: D

D ofcourse

upvoted 1 times

Titini 10 months ago

Selected Answer: D

D is correct

upvoted 1 times

jarz 1 year ago

Selected Answer: D

It does not offer protection in environments where IPv6 traffic is tunnelled.

upvoted 4 times

Question #354 Topic 1

A network administrator is troubleshooting a failed AAA login issue on a Cisco Catalyst c3560 switch. When the network administrator tries to log in with SSH using TACACS+ username and password credentials, the switch is no longer authenticating and is failing back to the local account. Which action resolves this issue?

A. Configure ip tacacs-server source-interface GigabitEthernet 1/1.

B. Configure ip tacacs source-ip 192.168.100.55.

C. Configure ip tacacs source-interface GigabitEthernet 1/1.

D. Configure ip tacacs-server source-ip 192.168.100.55.

Correct Answer: *B*

⊟ **DUBC89x** Highly Voted 1 year ago

Selected Answer: C

```
R1(config)#ip tacacs ?
source-interface Specify interface for source address in TACACS packets
```

upvoted 9 times

⊟ **Chiaretta** Most Recent 5 months, 1 week ago

Selected Answer: C

C is correct but i dont understand what it relate with question!!!

upvoted 3 times

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: C

option C is correct:

```
R2(config)#ip tacacs so
R2(config)#ip tacacs source-interface ?
Async Async interface
Auto-Template Auto-Template interface
BVI Bridge-Group Virtual Interface
CDMA-Ix CDMA Ix interface
CTunnel CTunnel interface
Dialer Dialer interface
FastEthernet FastEthernet IEEE 802.3
Lex Lex interface
Loopback Loopback interface
MFR Multilink Frame Relay bundle interface
Multilink Multilink-group interface
Null Null interface
Port-channel Ethernet Channel of interfaces
SSLVPN-VIF SSLVPN Virtual Interface
Tunnel Tunnel interface
Vif PGM Multicast Host interface
Virtual-PPP Virtual PPP interface
Virtual-Template Virtual Template interface
Virtual-TokenRing Virtual TokenRing
vmi Virtual Multipoint Interface

R2(config)#ip tacacs source-interface
```

upvoted 1 times

⊟ **HungarianDish** 7 months, 1 week ago

Selected Answer: C

The only valid command is "C" -> "ip tacacs source-interface".

Still, I do not think that I would solve the issue.

upvoted 2 times

⊟ **HungarianDish** 6 months, 3 weeks ago

This command does not solve the authentication issue. It only helps to troubleshoot, and makes log analyzes easier.
https://www.oreilly.com/library/view/cisco-ios-cookbook/0596527225/ch04s08.html

upvoted 1 times

⊟ **forccnp** 10 months ago

Selected Answer: C

C is the correct answer

upvoted 2 times

⊟ **forccnp** 10 months ago

```
Switch(config)#ip tacacs ?
source-interface Specify interface for source address in TACACS packets
```

upvoted 1 times

Question #355 *Topic 1*

Which two solutions are used to overcome a flapping link that causes a frequent label binding exchange between MPLS routers? (Choose two.)

A. Increase input queue on links to protect the session.

B. Increase a hold-timer to protect the session.

C. Increase a session delay to protect the session.

D. Create link dampening on links to protect the session.

E. Create targeted hellos to protect the session.

Correct Answer: *CD*

DUBC89x Highly Voted 1 year ago

Selected Answer: DE

To avoid having to rebuild the LDP session altogether, you can protect it. When the LDP session between two directly connected LSRs is protected, a targeted LDP session is built between the two LSRs. When the directly connected link does go down between the two LSRs, the targeted LDP session is kept up as long as an alternative path exists between the two LSRs.
For the protection to work, you need to enable it on both the LSRs. If this is not possible, you can enable it on one LSR, and the other LSR can accept the targeted LDP Hellos by configuring the command mpls ldp discovery targeted-hello accept.
Reference: https://www.ccexpert.us/mpls-network/mpls-ldp-session-protection.htmlOr from the referenceat
https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/5eU6DfQV/TECMPL-3201.pdfTroubleshooting LDP IssuesProblem:
I. When a link flaps (for a short time),
...S olution:
+ When LDP session supported by link hello is setup, create a targeted hello to protect the session.

upvoted 6 times

jansan55 Most Recent 2 months, 3 weeks ago

Selected Answer: BE

I agree with HungarianDish explanation. B and E try to keep the MPLS session alive.

upvoted 2 times

inteldarvid 4 months, 3 weeks ago

Selected Answer: DE

D and E:
To avoid having to rebuild the LDP session altogether, you can protect it. When the LDP session between two directly connected LSRs is protected, a targeted LDP session is built between the two LSRs. When the directly connected link does go down between the two LSRs, the targeted LDP session is kept up as long as an alternative path exists between the two LSRs. For the protection to work, you need to enable it on both the LSRs. If this is not possible, you can enable it on one LSR, and the other LSR can accept the targeted LDP Hellos by configuring the command mpls ldp discovery targeted-hello accept.
Reference:
https://www.ccexpert.us/mpls-network/mpls-ldp-session-protection.html
https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/5eU6DfQV/TECMPL-3201.pdf
Troubleshooting LDP Issues
Problem:
I. When a link flaps (for a short time),
...
Solution:
+ When LDP session supported by link hello is setup, create a targeted hello to protect the session.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: DE

D and E corerct:

team look the reference:
https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/5eU6DfQV/TECMPL-3201.pdf

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: BE

For me these are the closest:
B) Increase a hold-timer to protect the session.
E) Create targeted hellos to protect the session.

E) already explained by others.
https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/5eU6DfQV/TECMPL-3201.pdf
When link is down, the targeted hello remains through other path and keeps the LDP session up.

upvoted 4 times

HungarianDish 7 months, 1 week ago

B) Please, see:
https://community.cisco.com/t5/mpls/ldp-session-flap-on-a-7600-router/td-p/1926870
- increase the hold queue size on the interface
- also recommend you to configure mpls ldp sync

https://community.cisco.com/t5/mpls/bfd-with-ldp/td-p/2264164

In short, you could save your traffic from a flapping link by using a holddown timer.
Alternatively, you could use IP dampening on the interface which will give it a penalty on every successive flap and keep it down.
But, this will keep the interface completely down rather that the LDP-IGP Sync funda.
...
The holddown timer is used to ensure that IGP waits for LDP to be up for the specified interval before it stops advertizing the maximum metric and makes the route usable.

upvoted 2 times

Mad_Scorpion 10 months, 3 weeks ago

Selected Answer: DE

To avoid having to rebuild the LDP session altogether, you can protect it. When the LDP session between two directly connected LSRs is protected, a targeted LDP session is built between the two LSRs. When the directly connected link does go down between the two LSRs, the targeted LDP session is kept up as long as an alternative path exists between the two LSRs.
For the protection to work, you need to enable it on both the LSRs. If this is not possible, you can enable it on one LSR, and the other LSR can accept the targeted LDP Hellos by configuring the command mpls ldp discovery targeted-hello accept.

upvoted 1 times

Question #356

Topic 1

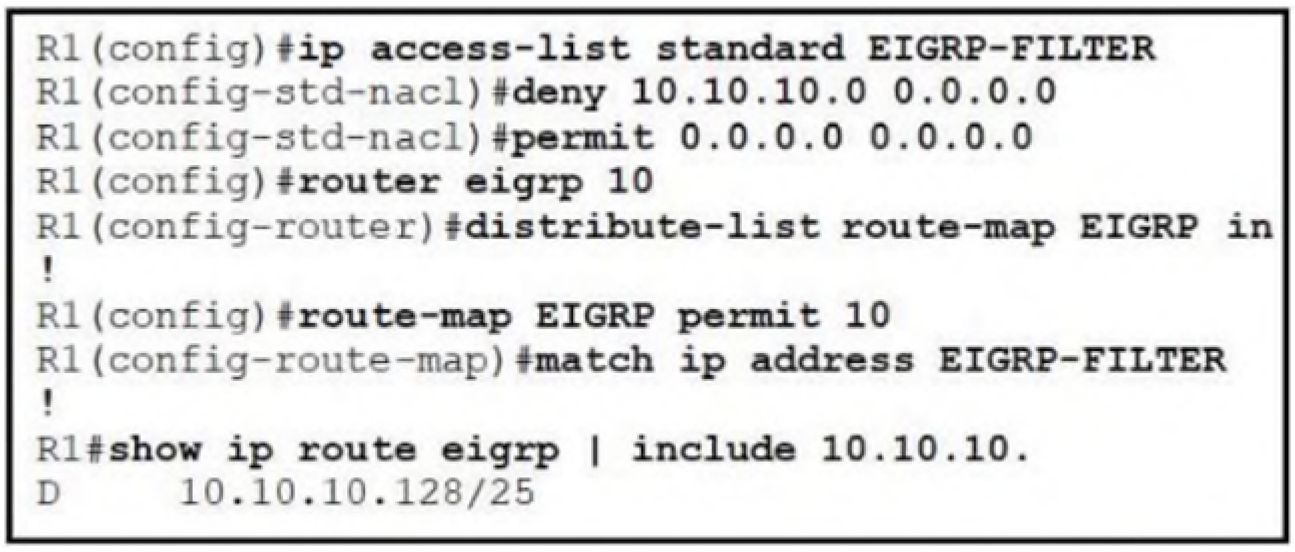

```
R1(config)#ip access-list standard EIGRP-FILTER
R1(config-std-nacl)#deny 10.10.10.0 0.0.0.0
R1(config-std-nacl)#permit 0.0.0.0 0.0.0.0
R1(config)#router eigrp 10
R1(config-router)#distribute-list route-map EIGRP in
!
R1(config)#route-map EIGRP permit 10
R1(config-route-map)#match ip address EIGRP-FILTER
!
R1#show ip route eigrp | include 10.10.10.
D      10.10.10.128/25
```

Refer to the exhibit. An engineer must filter EIGRP updates that are received to block all 10.10.10.0/24 prefixes. The engineer tests the distribute list and finds one associated prefix. Which action resolves the issue?

A. There is a permit in the ACL that allows this prefix into EIGRP. The ACL should be modified to deny 10.10.10.0 255.255.255.0.

B. There is a permit in the ACL that allows this prefix into EIGRP. The ACL should be modified to deny 10.10.10.0 0.0.0.255.

C. There is a permit in the route map that allows this prefix. A deny 20 statement is required with a match condition to match a new ACL that denies all prefixes.

D. There is a permit in the route map that allows this prefix. A deny 20 statement is required with no match condition to block the prefix.

Correct Answer: *B*

HungarianDish 7 months, 1 week ago

If they aim to filter 10.10.10.0/24, then I would expect to see something like this:
ip access-list standard EIGRP-FILTER
permit 10.10.10.0 0.0.0.255

route-map EIGRP deny 10
match ip address EIGRP-FILTER
route-map EIGRP permit 20

upvoted 3 times

DeWalt95 3 weeks, 5 days ago

Agree that B is closest to a coherent answer.

upvoted 1 times

inteldarvid 5 months ago

you have a reason, this question is wrong. But I need choice on aswer. I thinsg option B.

upvoted 1 times

sasasan12345 9 months, 1 week ago

Selected Answer: B

ACL uses wildcard mask.

upvoted 2 times

Typovy 9 months, 1 week ago

Selected Answer: A

This is standard ACL, A is the correct answer

upvoted 1 times

Typovy 8 months, 2 weeks ago

My bad even standard acl use wildcard

upvoted 1 times

Xerath 9 months ago

ACLs use wildcard masks, not subnet masks.

upvoted 1 times

Question #357 *Topic 1*

A network engineer must configure a DMVPN network so that a spoke establishes a direct path to another spoke if the two must send traffic to each other. A spoke must send traffic directly to the hub if required. Which configuration meets this requirement?

A. At the hub router:

```
interface tunnel10
ip nhrp nhs dynamic multipoint
ip nhrp nhs shortcut
tunnel mode gre multicast
```

On the spokes router:

```
interface tunnel10
ip nhrp nhs multicast dynamic
ip nhrp nhs redirect
tunnel mode gre multicast
```

B. At the hub router:

```
interface tunnel10
ip nhrp map dynamic multipoint
ip nhrp redirect
tunnel mode gre multicast
```

On the spokes router:

```
interface tunnel10
ip nhrp map multicast dynamic
ip nhrp shortcut
tunnel mode gre multicast
```

C. At the hub router:

```
interface tunnel10
ip nhrp nhs multicast dynamic
ip nhrp nhs shortcut
tunnel mode gre multipoint
```

On the spokes router:

```
interface tunnel10
ip nhrp nhs multicast dynamic
ip nhrp nhs redirect
tunnel mode gre multipoint
```

D. At the hub router:

```
interface tunnel10
ip nhrp map multicast dynamic
ip nhrp redirect
tunnel mode gre multipoint
```

On the spokes router:

```
interface tunnel10
ip nhrp map multicast dynamic
ip nhrp shortcut
tunnel mode gre multipoint
```

Correct Answer: *D*

Xerath 9 months, 3 weeks ago

Selected Answer: D

The given answer is correct.

upvoted 1 times

forccnp 10 months ago

Selected Answer: D

D is the corrrect
ez

upvoted 1 times

DUBC89x 1 year ago

Selected Answer: D

Verified it with lab with Network Lessons.

upvoted 2 times

Question #358 *Topic 1*

The network administrator configured R1 to authenticate Telnet connections based on Cisco ISE using TACACS+. ISE has been configured with an IP address of 192.168.1.5 and with a network device pointing toward R1 (192.168.1.1) with a shared secret password of Cisco123.

The administrator has configured this on R1:

```
aaa new-model
!
tacacs server ISE1
address ipv4 192.168.1.5
key Cisco123
!
aaa group server tacacs+ TAC-SERV
server name ISE1
!
aaa authentication login telnet group TAC-SERV
```

The network administrator cannot authenticate to R1 based on ISE. Which configuration fixes the issue?

A. line vty 0 4
login authentication TAC-SERV

B. tacacs-server host 192.168.1.5 key Cisco123

C. ip tacacs-server host 192.168.1.5 key Cisco123

D. line vty 0 4
login authentication telnet

Correct Answer: *A*

SAMAKEMM 2 months, 1 week ago

Selected Answer: D

D is correct

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: D

yes optiion D

upvoted 1 times

tapri 10 months, 3 weeks ago

Selected Answer: D

D is correct

upvoted 3 times

ellen_AA 11 months, 2 weeks ago

aaa authentication methode called "telnet" was created, but not used under the vty lines.

upvoted 3 times

mitosenoriko 11 months, 2 weeks ago

i think A

```
  (config)# line [ console | vty | tty | aux ] number number
  (config-line)# login authentication [ default | list-name ]
```

upvoted 2 times

Zizu007 11 months, 2 weeks ago

Selected Answer: D

```
aaa new-model
!
tacacs server ISE1
address ipv4 192.168.1.5
key Cisco123
!
```

```
aaa group server tacacs+ TAC-SERV
server name ISE1
!
aaa authentication login telnet group TAC-SERV
!
R5(config)#line vty 0 4
R5(config-line)#login authentication TAC-SERV
AAA: Warning authentication list "TAC-SERV" is not defined for LOGIN.

R5(config-line)#login authentication telnet
R5(config-line)#end
R5#
```

upvoted 4 times

JKStinn 1 year ago

Selected Answer: D

https://www.omnisecu.com/ccna-security/cisco-router-switch-aaa-login-authentication-configuration-using-tacacs+-and-radius-protocols-through-commands.php

upvoted 4 times

Question #359 *Topic 1*

The network administrator must configure R1 to authenticate Telnet connections based on Cisco ISE using RADIUS. ISE has been configured with an IP address of 192.168.1.5 and with a network device pointing toward R1 (192.168.1.1) with a shared secret password of Cisco123. The administrator has configured this on R1:

```
aaa new-model
!
radius server ISE1
address ipv4 192.168.1.5
key Cisco123
!
aaa group server tacacs+ RAD-SERV
server name ISE1
!
aaa authentication login default group RAD-SERV
```

The network administrator cannot authenticate to access R1 based on ISE. Which set of configurations fixes the issue?

A. line vty 0 4
login authentication RAD-SERV

B. aaa group server tacacs+ ISE1
server name RAD-SERV

C. aaa group server radius RAD-SERV
server name ISE1

D. line vty 0 4
login authentication default

Correct Answer: *A*

Zizu007 Highly Voted 11 months, 3 weeks ago

Selected Answer: C

"aaa group server tacacs+" is incorrect in this case!

```
aaa new-model
aaa authentication login default group radius
!
aaa group server radius RAD
server name SERV_RADIUS
!
radius server SERV_RADIUS
address ipv4 <RADIUS-ADDRESS>
key <string>
!
```

upvoted 5 times

inteldarvid Most Recent 5 months, 1 week ago

Selected Answer: C

option C 100 %%% correct

upvoted 2 times

ellen_AA 11 months, 2 weeks ago

Selected Answer: C

Correct answer is C:
Router is using default method: #aaa authentication login default group RAD-SERV
that calls a group RAD-SERV (which is a tacas+ in this scenario), but the question asks about a radius server config, not tacacs+. So we need to configure a radius group named RAD-SERV.

upvoted 4 times

Question #360 *Topic 1*

Which IPv6 first-hop security feature helps to minimize denial of service attacks?

A. IPv6 Router Advertisement Guard

B. IPv6 Destination Guard

C. DHCPv6 Guard

D. IPv6 MAC address filtering

Correct Answer: *B*

yeyuno 8 months, 2 weeks ago

IPv6 - Destination Guard
The Destination Guard feature helps in minimizing denial-of-service (DoS) attacks. It performs address resolutions only for those addresses that are active on the link, and requires the FHS binding table to be populated with the help of the IPv6 snooping feature.

The feature enables the filtering of IPv6 traffic based on the destination address, and blocks the NDP resolution for destination addresses that are not found in the binding table. By default, the policy drops traffic coming for an unknown destination.

upvoted 1 times

Titini 10 months, 1 week ago

Selected Answer: B

https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/15S/configuration/guide/7600_15_0s_book/IPv6_Security.html

upvoted 2 times

kaisehhop 11 months ago

Selected Answer: B

B is Correct
https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/15S/configuration/guide/7600_15_0s_book/IPv6_Security.html

upvoted 3 times

mitosenoriko 11 months, 2 weeks ago

C is correct.
I checked cisco document.

upvoted 2 times

> **onkel_andi** 11 months ago
>
> B is correct, check the Cisco document, LOL
>
> upvoted 3 times
>
> > **forccnp** 10 months ago
> >
> > :DDD
> > B is correct
> >
> > upvoted 2 times

Question #361 *Topic 1*

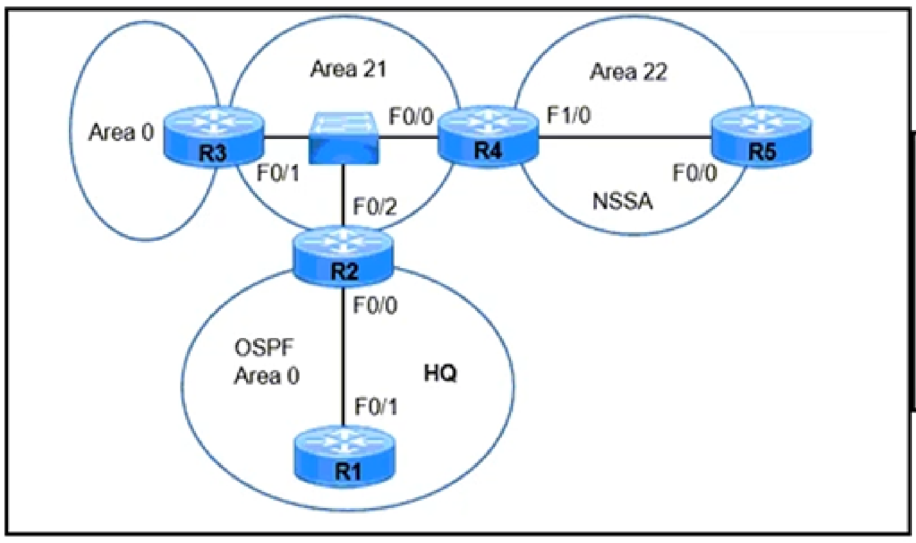

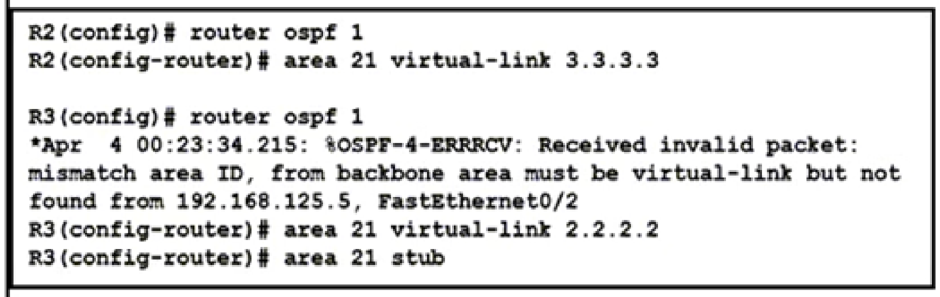

```
R2(config)# router ospf 1
R2(config-router)# area 21 virtual-link 3.3.3.3

R3(config)# router ospf 1
*Apr  4 00:23:34.215: %OSPF-4-ERRRCV: Received invalid packet:
mismatch area ID, from backbone area must be virtual-link but not
found from 192.168.125.5, FastEthernet0/2
R3(config-router)# area 21 virtual-link 2.2.2.2
R3(config-router)# area 21 stub
```

Refer to the exhibit. A network engineer is troubleshooting a failed link between R2 and R3. No traffic loss is reported from router R5 to HQ. Which command fixes the separated backbone?

A. R3(config-router)#area 21 virtual-link 192.168.125.5

B. R2(config-router)#area 21 virtual-link 192.168.125.5

C. R3(config-router)#no area 21 stub

D. R2(config-router)#no area 21 stub

Correct Answer: *B*

everfly Highly Voted 1 year ago

Selected Answer: C

R3(config-router)#no area 21 stub

upvoted 5 times

inteldarvid Most Recent 5 months, 1 week ago

Selected Answer: C

yes, option C correct

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: C

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/47866-ospfdb7.html

The area through which you configure the virtual link, known as a transit area , must have full routing information.
The transit area cannot be a stub area.

upvoted 2 times

steve_lee 7 months, 1 week ago

Selected Answer: C

I vote C.
Virtual Link can't formed in any stub area.

upvoted 3 times

Titini 10 months ago

Selected Answer: C

The error log is printed prior configuring the virtual link and the stub area. The log actually confirms the fact that r2 tried to establish a virtual link with r3. The 192.168.125.5 is the source IP of the packet that triggered the error and belongs to R2. it is not necessarily the same as the Router ID. Anyway, if the virtual link endpoints are in the stub area, then connectivity between the two routers could be lost. To allow a virtual link to exist through a stub area you should use "virtual-link stub".

upvoted 2 times

Question #362

Topic 1

A CoPP policy is applied for receiving SSH traffic from the WAN interface on a Cisco ISR4321 router. However, the SSH response from the router is abnormal and stuck during the high link utilization. The problem is identified as SSH traffic does not match in the ACL. Which action resolves the issue?

A. Apply CoPP on the control plane interface.

B. Apply CoPP on the WAN interface inbound direction.

C. Rate-limit SSH traffic to ensure dedicated bandwidth.

D. Increase the IP precedence value of SSH traffic to 6.

Correct Answer: *A*

[Removed] 4 months ago

Selected Answer: A

There is only one control plane interface.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: A

option A correct:
The problem is “SSH traffic does not match in the ACL” and “CoPP policy is applied for receiving SSH traffic from the WAN interface” so we should apply CoPP on the control plane interface instead.

upvoted 1 times

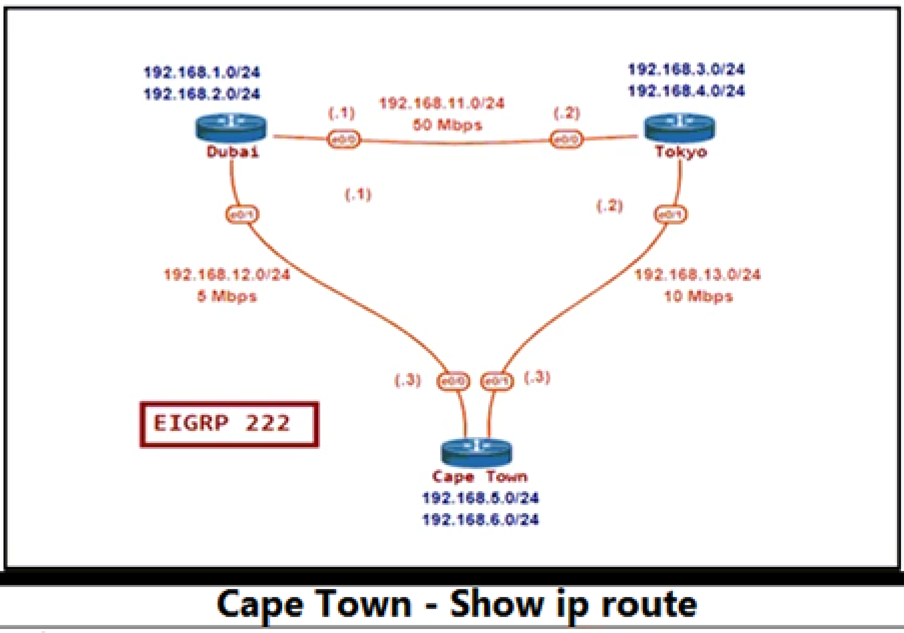

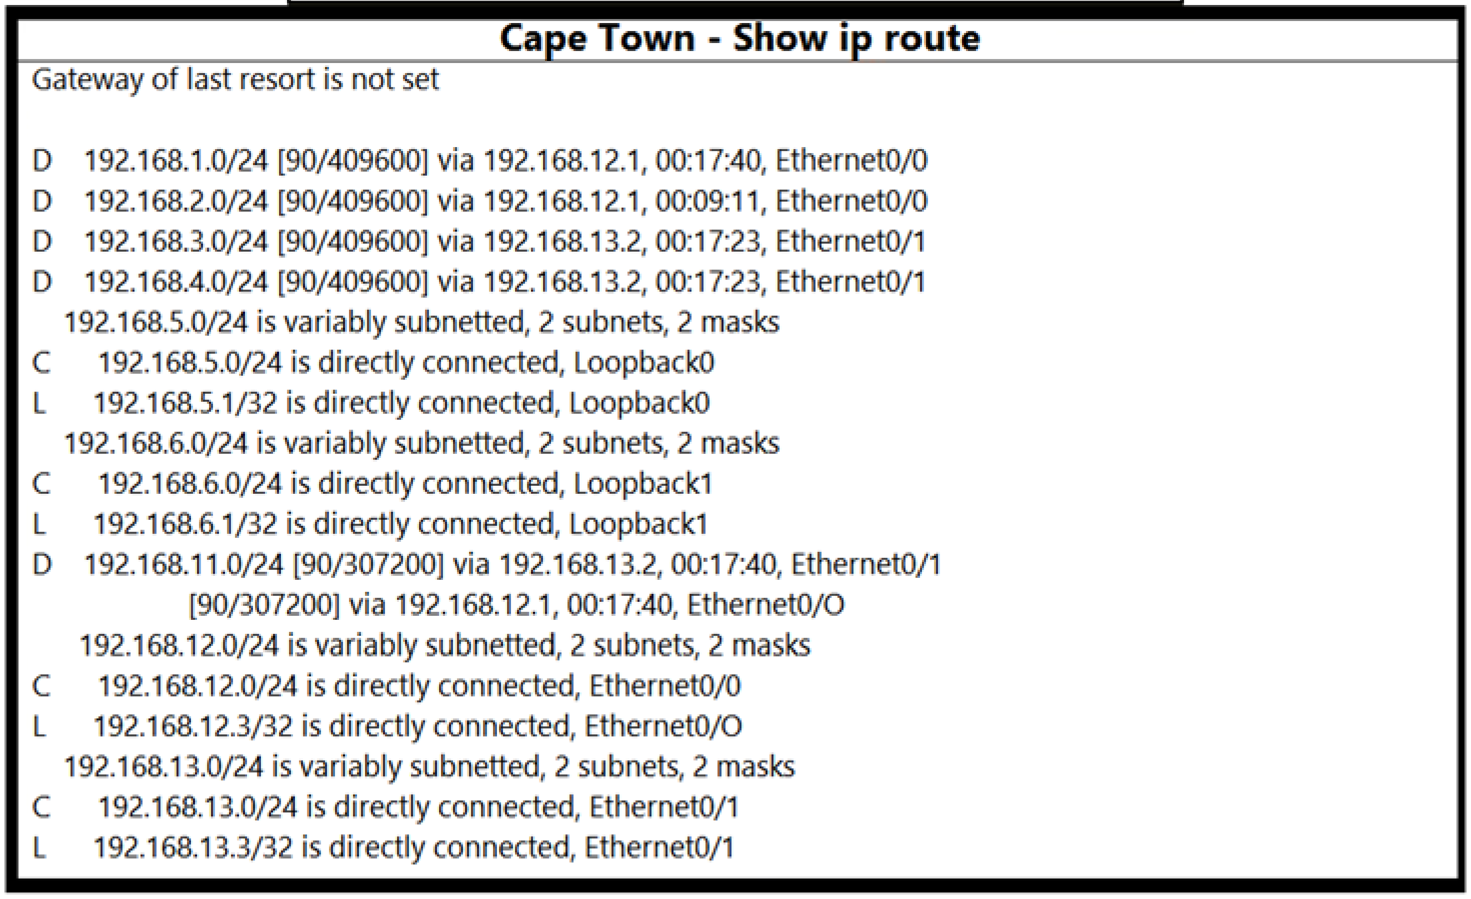

Cape Town - Show ip route

```
Gateway of last resort is not set

D   192.168.1.0/24 [90/409600] via 192.168.12.1, 00:17:40, Ethernet0/0
D   192.168.2.0/24 [90/409600] via 192.168.12.1, 00:09:11, Ethernet0/0
D   192.168.3.0/24 [90/409600] via 192.168.13.2, 00:17:23, Ethernet0/1
D   192.168.4.0/24 [90/409600] via 192.168.13.2, 00:17:23, Ethernet0/1
      192.168.5.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.5.0/24 is directly connected, Loopback0
L      192.168.5.1/32 is directly connected, Loopback0
      192.168.6.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.6.0/24 is directly connected, Loopback1
L      192.168.6.1/32 is directly connected, Loopback1
D   192.168.11.0/24 [90/307200] via 192.168.13.2, 00:17:40, Ethernet0/1
                  [90/307200] via 192.168.12.1, 00:17:40, Ethernet0/O
       192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.12.0/24 is directly connected, Ethernet0/0
L      192.168.12.3/32 is directly connected, Ethernet0/O
      192.168.13.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.13.0/24 is directly connected, Ethernet0/1
L      192.168.13.3/32 is directly connected, Ethernet0/1
```

Refer to the exhibit. The network administrator must configure Cape Town to reach Dubai via Tokyo based on the speeds provided by the service provider. It was noticed that Cape Town is reaching Dubai directly and failed to meet the requirement. Which configuration fixes the issue?

A. CapeTown -

router eigrp 100
variance 2

B. CapeTown -

interface E 0/0
bandwidth 5000
interface E 0/1
bandwidth 10000

C. CapeTown -

interface E 0/0
bandwidth 5000
interface E 0/1
bandwidth 10000

Dubai -

interface E 0/0

```
bandwidth 50000
interface E 0/1
bandwidth 5000
```

Tokyo -

```
interface E 0/0
bandwidth 50000
interface E 0/1
bandwidth 10000
```

D. Dubai -

```
router eigrp 100
variance 2
```

Correct Answer: *B*

Gedson 4 months, 4 weeks ago

Selected Answer: B

The network administrator must configure Cape Town to reach Dubai via Tokyo based on the speeds provided by the service provider. It was noticed that Cape Town is reaching Dubai directly and failed to meet the requirement. Which configuration fixes the issue?

upvoted 1 times

Gedson 4 months, 4 weeks ago

Bw represents the slowest link in the path. Link speed is colected form the confgured interface bandwidth.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

yes anwser correct is "C", all path

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: C

https://community.cisco.com/t5/switching/changing-interface-bandwidth-affects-eigrp-availability/td-p/2879429

upvoted 1 times

shoo83 11 months, 1 week ago

Selected Answer: C

to modify cost, we must modify bandwith all path

upvoted 4 times

ellen_AA 11 months, 2 weeks ago

If B is correct, why not C?

upvoted 3 times

Almylle 5 months, 4 weeks ago

Probably because u only need to manipulate the metric of the Cape town to tokyo and then to dubai, so assume that tokyo and dubai still works with their default metrics

upvoted 1 times

Question #364 Topic 1

DRAG DROP

-

Drag and drop the ICMPv6 neighbor discovery messages from the left onto the correct packet types on the right.

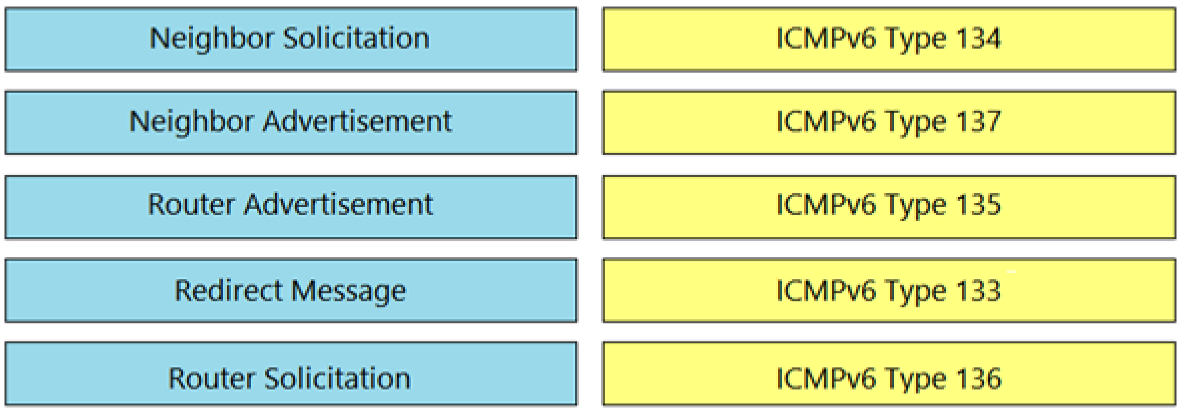

Neighbor Solicitation	ICMPv6 Type 134
Neighbor Advertisement	ICMPv6 Type 137
Router Advertisement	ICMPv6 Type 135
Redirect Message	ICMPv6 Type 133
Router Solicitation	ICMPv6 Type 136

Correct Answer:

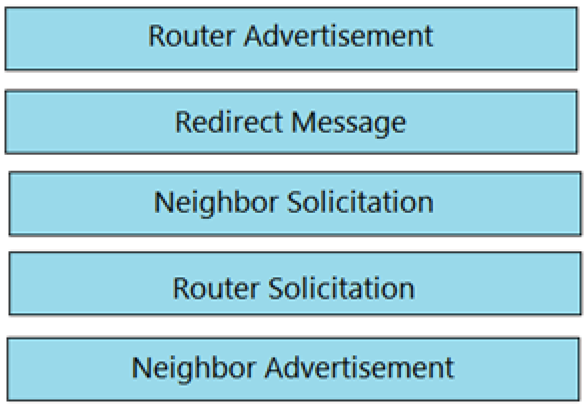

Router Advertisement

Redirect Message

Neighbor Solicitation

Router Solicitation

Neighbor Advertisement

abd123 Highly Voted 10 months, 2 weeks ago

really cisco !! , memorizing Q for CCNP level

upvoted 12 times

forccnp Highly Voted 10 months ago

RS - 133
RA - 134
NS - 135
NA - 136
RM - 137

upvoted 7 times

inteldarvid Most Recent 5 months, 1 week ago

I found a pattern to learn better, from least to greatest (133,134,135,136,137), I start with router solicitation, Router Ad. Nei Request, Neig Ad, redirect

upvoted 1 times

inteldarvid 5 months, 1 week ago

WTF really CISCO, really???????????????.

upvoted 2 times

kaisehhop 11 months ago

The answer shown in the image is correct
https://www.cisco.com/c/en/us/support/docs/ip/routing-information-protocol-rip/22974-icmpv6codes.html

upvoted 1 times

Zizu007 11 months, 3 weeks ago

correct!

Type 133 - Router Solicitation
Type 134 - Router Advertisement
Type 135 - Neighbor Solicitation
Type 136 - Neighbor Advertisement
Type 137 - Redirect Message

upvoted 3 times

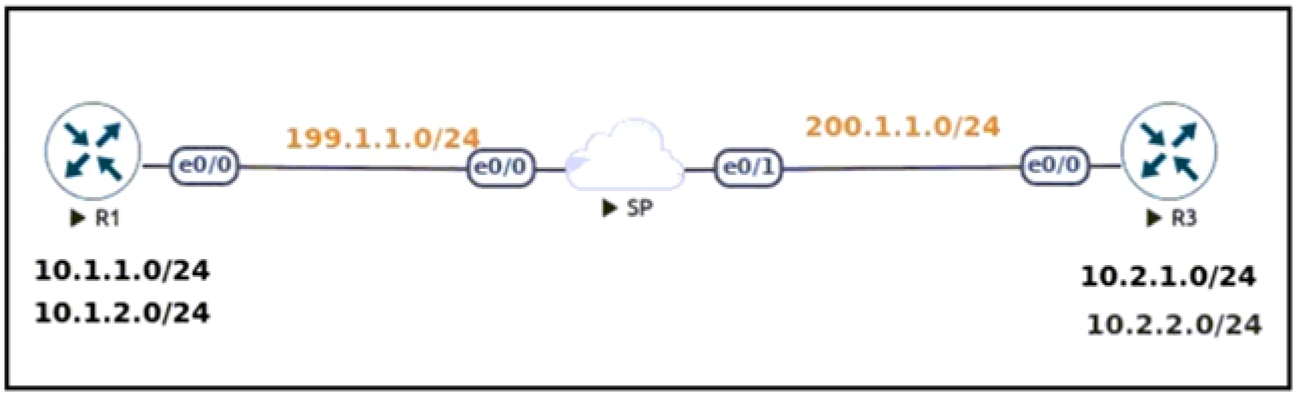

Refer to the exhibit. An engineer must configure a LAN-to-LAN IPsec VPN between R1 and the remote router. Which IPsec Phase 1 configuration must the engineer use for the local router?

A. crypto isakmp policy 5
authentication pre-share
encryption 3des
hash sha
group 2
!
crypto isakmp key cisco123 address 200.1.1.3

B. crypto isakmp policy 5
authentication pre-share
encryption 3des
hash md5
group 2
!
crypto isakmp key cisco123! address 199.1.1.1

C. crypto isakmp policy 5
authentication pre-share
encryption 3des
hash md5
group 2
!
crypto isakmp key cisco123 address 199.1.1.1

D. crypto isakmp policy 5
authentication pre-share
encryption 3des
hash md5
group 2
!
crypto isakmp key cisco123 address 200.1.1.3

Correct Answer: *D*

ZamanR 4 days, 22 hours ago

A is correct answer
Explanation

In the "crypto isakmp key … address " command, the address must be of the IP address of the other end (which is 200.1.1.3 in this case) so Option A and Option B are correct. The difference between these two options are in the hash SHA or MD5 method but both of them can be used although SHA is better than MD5 so we choose Option A the best answer.

Note: Cisco no longer recommends using 3DES, MD5 and DH groups 1, 2 and 5.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_imgmt/configuration/xe-16-

5/sec-ipsec-management-xe-16-5-book/sec-ipsec-usability-enhance.html

upvoted 1 times

HarwinderSekhon 4 months, 1 week ago

Selected Answer: A

A and D both are correct but remember to smoke before exam :P

```
vIOS(config-isakmp)#hash ?
md5 Message Digest 5
sha Secure Hash Standard
sha256 Secure Hash Standard 2 (256 bit)
sha384 Secure Hash Standard 2 (384 bit)
sha512 Secure Hash Standard 2 (512 bit)
```

upvoted 1 times

guy276465281819372 4 months, 1 week ago

Selected Answer: D

D would be fast and simple, A more secure. no way to choose.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: A

100 % "A", because sha is more safe than md5

upvoted 2 times

pepgua 5 months, 2 weeks ago

Selected Answer: D

SHA or SHA1 ? A doesn't look correct?

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: A

Agree with others.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_ikevpn/configuration/xe-16-5/sec-ike-for-ipsec-vpns-xe-16-5-book/sec-key-exch-ipsec.html

Cisco no longer recommends using DES, 3DES, MD5 (including HMAC variant), and Diffie-Hellman (DH) groups 1, 2 and 5; instead, you should use AES, SHA-256 and DH Groups 14 or higher.

upvoted 1 times

azzawim 8 months, 4 weeks ago

Selected Answer: A

correct answer A

upvoted 2 times

sayed_2908 10 months, 3 weeks ago

Selected Answer: A

A & D is correct but SHA is safer than MD5. hence I choose A.

upvoted 3 times

dq28 11 months, 2 weeks ago

I agree ... I think there is something missing in the question. Sha is safer, md5 is faster. Which one should be chosen? Only Cisco (if any) knows.

upvoted 4 times

mitosenoriko 11 months, 2 weeks ago

A and D is correct.
i dont select one.

upvoted 2 times

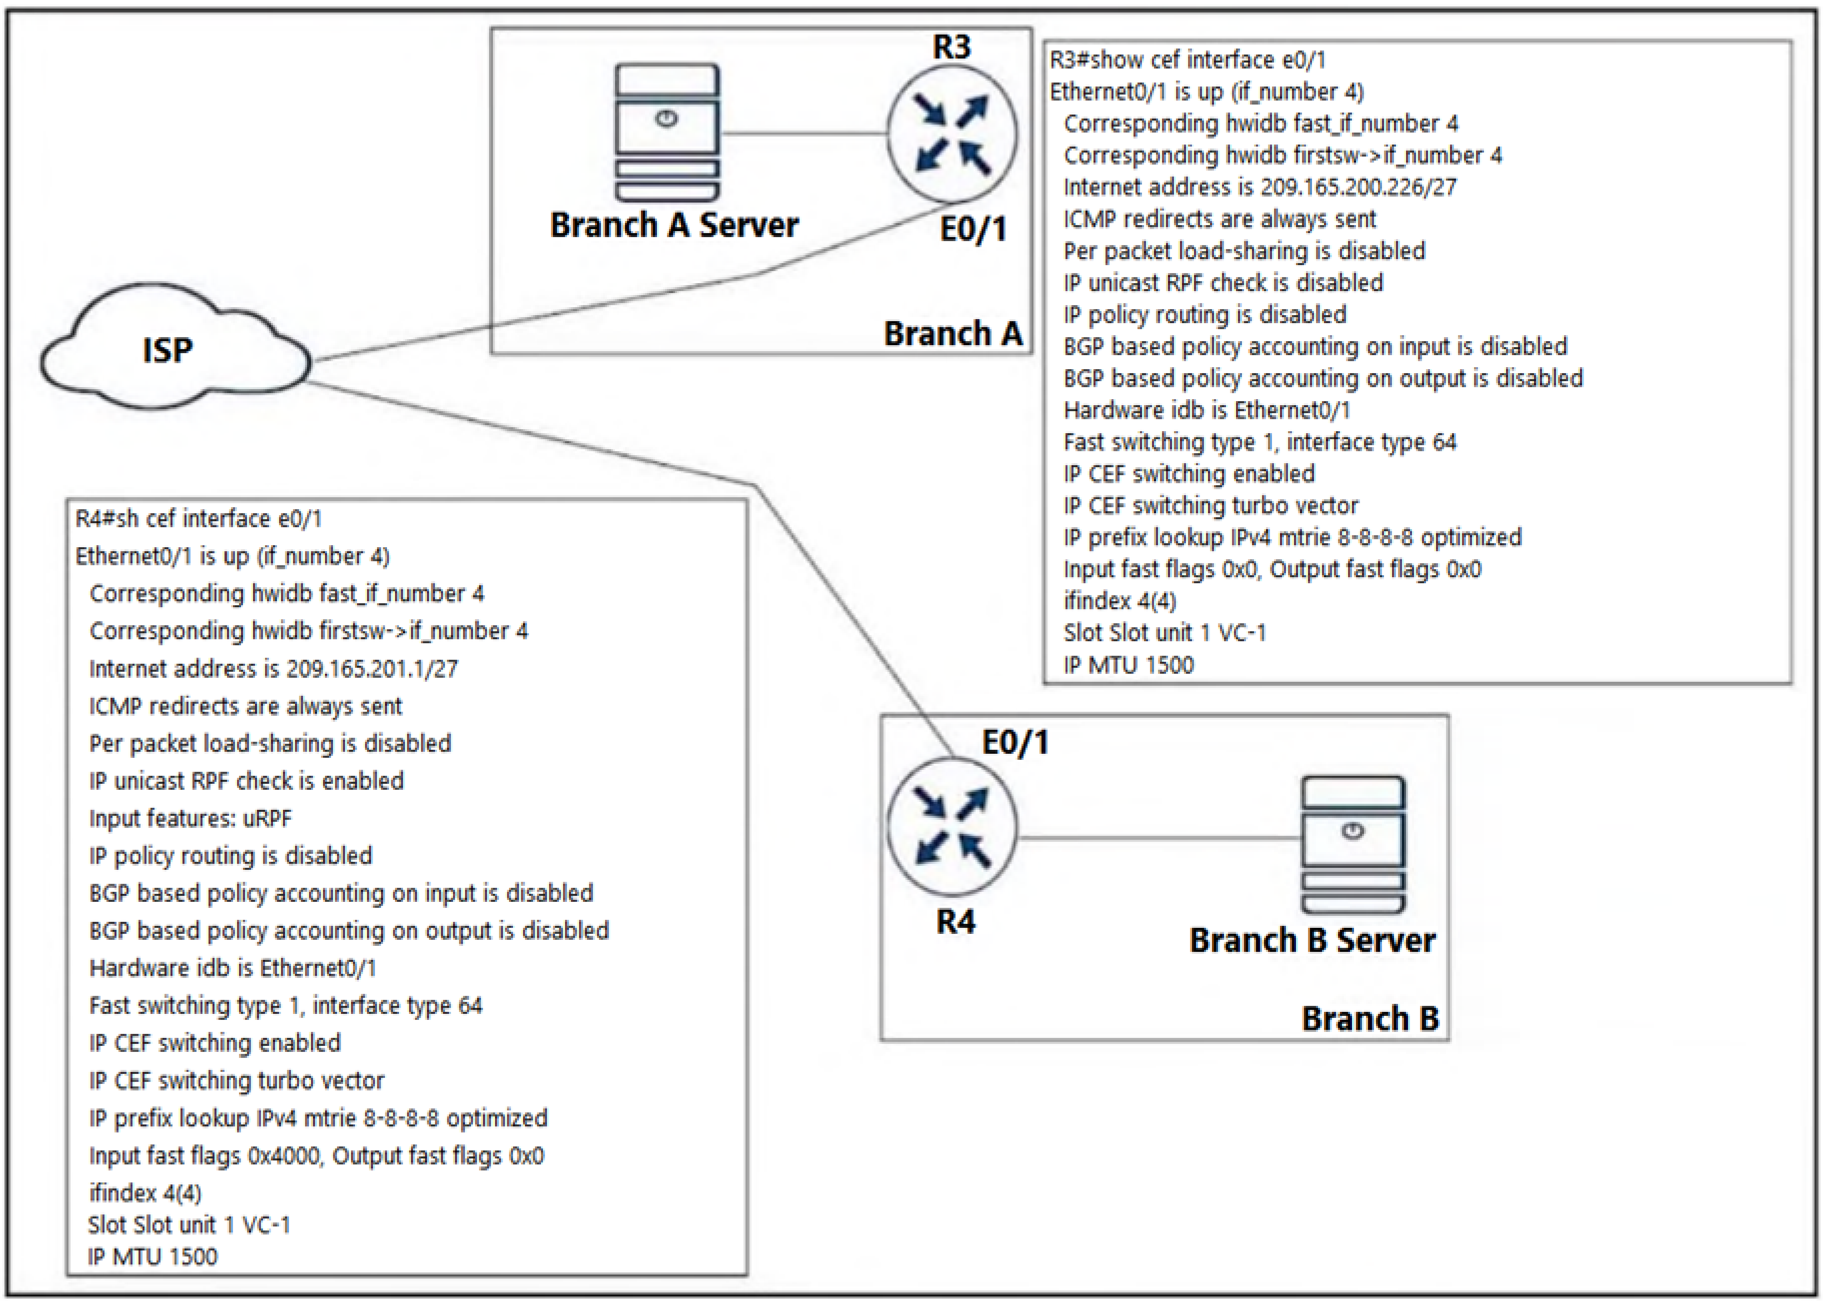

```
R3#show cef interface e0/1
Ethernet0/1 is up (if_number 4)
  Corresponding hwidb fast_if_number 4
  Corresponding hwidb firstsw->if_number 4
  Internet address is 209.165.200.226/27
  ICMP redirects are always sent
  Per packet load-sharing is disabled
  IP unicast RPF check is disabled
  IP policy routing is disabled
  BGP based policy accounting on input is disabled
  BGP based policy accounting on output is disabled
  Hardware idb is Ethernet0/1
  Fast switching type 1, interface type 64
  IP CEF switching enabled
  IP CEF switching turbo vector
  IP prefix lookup IPv4 mtrie 8-8-8-8 optimized
  Input fast flags 0x0, Output fast flags 0x0
  ifindex 4(4)
  Slot Slot unit 1 VC-1
  IP MTU 1500
```

```
R4#sh cef interface e0/1
Ethernet0/1 is up (if_number 4)
  Corresponding hwidb fast_if_number 4
  Corresponding hwidb firstsw->if_number 4
  Internet address is 209.165.201.1/27
  ICMP redirects are always sent
  Per packet load-sharing is disabled
  IP unicast RPF check is enabled
  Input features: uRPF
  IP policy routing is disabled
  BGP based policy accounting on input is disabled
  BGP based policy accounting on output is disabled
  Hardware idb is Ethernet0/1
  Fast switching type 1, interface type 64
  IP CEF switching enabled
  IP CEF switching turbo vector
  IP prefix lookup IPv4 mtrie 8-8-8-8 optimized
  Input fast flags 0x4000, Output fast flags 0x0
  ifindex 4(4)
  Slot Slot unit 1 VC-1
  IP MTU 1500
```

Refer to the exhibit. A shoe retail company implemented the uRPF solution for an antispoofing attack. A network engineer received the call that the branch A server is under an IP spoofing attack. Which configuration must be implemented to resolve the attack?

A. R4 -
interface ethernet0/1
ip verify unicast source reachable-via any allow-default allow-self-ping

B. R4 -
interface ethernet0/1
ip unicast RPF check reachable-via any allow-default allow-self-ping

C. R3 -
interface ethernet0/1
ip verify unicast source reachable-via any allow-default allow-self-ping

D. R3 -
interface ethernet0/1
ip unicast RPF check reachable-via any allow-default allow-self-ping

Correct Answer: *C*

⊟ **mitosenoriko** Highly Voted 11 months, 2 weeks ago

C is Correct
Router(config-if)# ip verify unicast source reachable-via any allow-self-ping

upvoted 5 times

⊟ **inteldarvid** Most Recent 5 months, 1 week ago

Selected Answer: C

option c: correct

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_urpf/configuration/xe-3s/sec-data-urpf-xe-3s-book/cfg-unicast-rpf.html

upvoted 1 times

Question #367 *Topic 1*

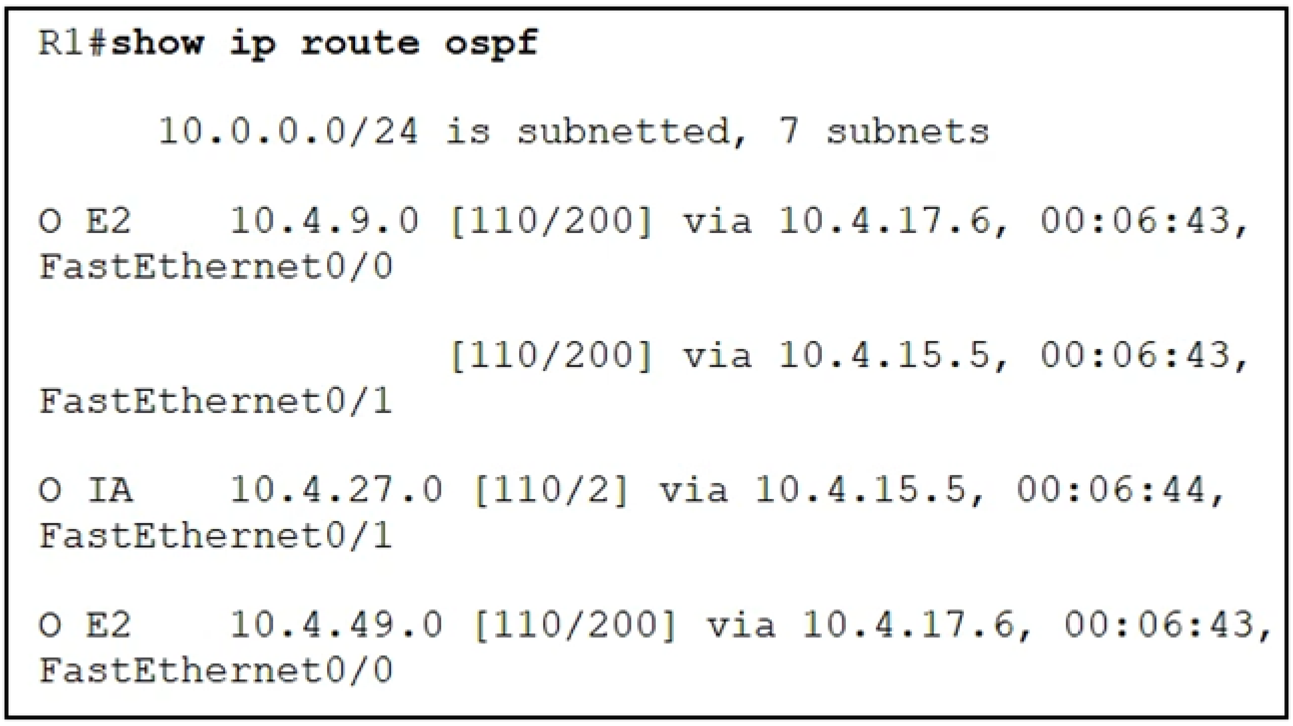

```
R1#show ip route ospf

     10.0.0.0/24 is subnetted, 7 subnets

O E2    10.4.9.0 [110/200] via 10.4.17.6, 00:06:43,
FastEthernet0/0

                 [110/200] via 10.4.15.5, 00:06:43,
FastEthernet0/1

O IA    10.4.27.0 [110/2] via 10.4.15.5, 00:06:44,
FastEthernet0/1

O E2    10.4.49.0 [110/200] via 10.4.17.6, 00:06:43,
FastEthernet0/0
```

Refer to the exhibit. An engineer configures two ASBRs, 10.4.17.6 and 10.4.15.5, in an OSPF network to redistribute routes from EIGRP. However, both ASBRs show the EIGRP routes as equal costs even though the next-hop router 10.4.17.6 is closer to R1. How should the network traffic to the EIGRP prefixes be sent via 10.4.17.6?

A. The administrative distance should be raised to 120 from the ASBR 10.4.17.6.

B. The redistributed prefixes should be advertised as Type 1.

C. The ASBR 10.4.17.6 should assign a tag to match and assign a lower metric on R1.

D. The administrative distance should be raised to 120 from the ASBR 10.4.15.5.

Correct Answer: *B*

Tester948 1 month ago

Selected Answer: B

Type 1 routes increment metric hop by hop, type 2 doesn't - a router that's 1 hop away will have the same metric as a router 30 hops away.

upvoted 2 times

Fenix7 3 months, 1 week ago

The answer is B. When there are multiple ASBB, you need to use E1.

* E1: Type O E1 external routes calculate the cost by adding the external cost to the internal cost of each link that the packet crosses. Use this type to avoid suboptimal routing when there are multiple ASBRs advertising an external route to the same AS.

* E2 (default): The external cost of O E2 packet routes is always the external cost only. Use this type if only one ASBR is advertising an external route to the AS.

upvoted 1 times

Pepoydex 3 months, 3 weeks ago

Selected Answer: D

I think the answer is d, because, the 2 route have the same distance....we want go fot 6.6

upvoted 1 times

HarwinderSekhon 4 months, 1 week ago

given ans is correct

upvoted 1 times

Question #368

Topic 1

Which component of MPLS VPNs is used to extend the IP address so that an engineer is able to identify to which VPN it belongs?

A. RT

B. RD

C. LDP

D. VPNv4 address family

Correct Answer: *B*

Stylar 5 months ago

VPNv4 Addressing is used to identify specific routes within a VPN. They are typically represented by an RD and an IPv4 prefix.
The combination of the RD and the IPv4 prefix forms a unique VPNv4 address for each route
The distribution of VPN routing information is controlled through the use of VPN route target communities

upvoted 2 times

HungarianDish 7 months, 1 week ago

Selected Answer: B

https://www.ciscolive.com/c/dam/r/ciscolive/apjc/docs/2019/pdf/BRKCRT-2601.pdf
https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2020/pdf/DGTL-BRKMPL-1100.pdf
https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/BRKMPL-1100.pdf

RD is prepended to each prefix to make routes unique.
VPNv4 address (96-bit address): 64-bit RD + 32 bit IPv4 address

upvoted 1 times

Question #369 *Topic 1*

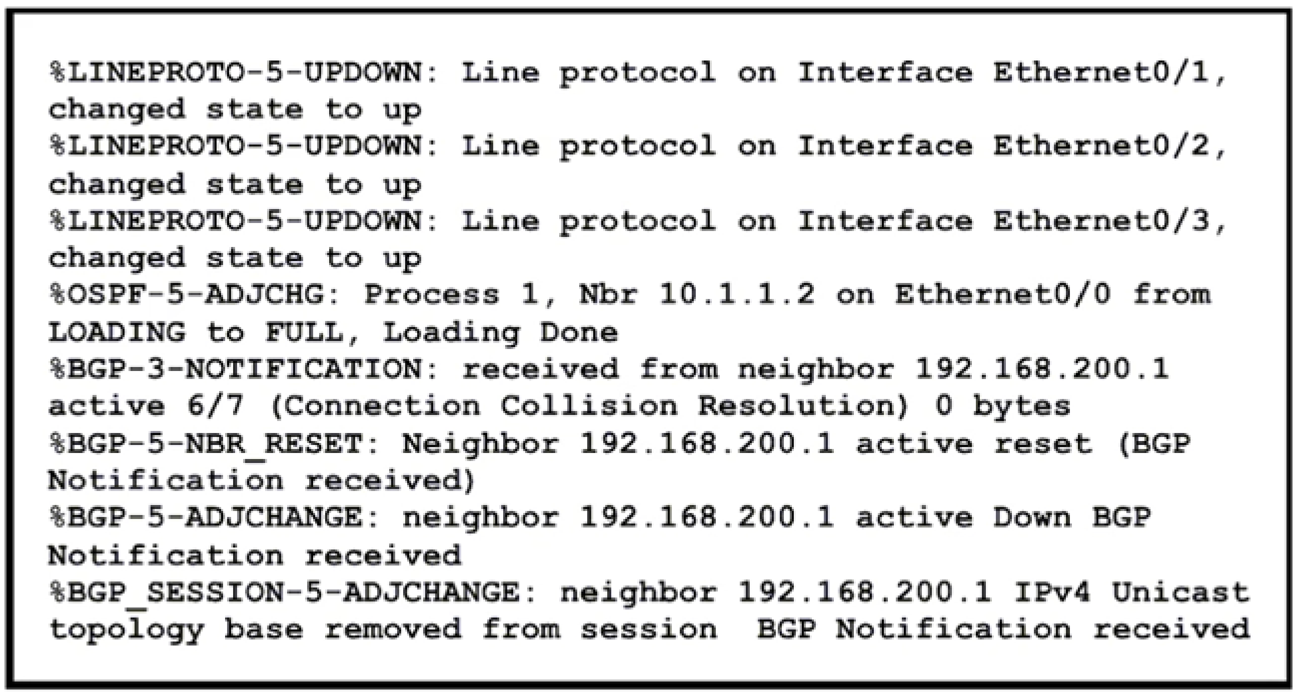

```
%LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0/1,
changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0/2,
changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0/3,
changed state to up
%OSPF-5-ADJCHG: Process 1, Nbr 10.1.1.2 on Ethernet0/0 from
LOADING to FULL, Loading Done
%BGP-3-NOTIFICATION: received from neighbor 192.168.200.1
active 6/7 (Connection Collision Resolution) 0 bytes
%BGP-5-NBR_RESET: Neighbor 192.168.200.1 active reset (BGP
Notification received)
%BGP-5-ADJCHANGE: neighbor 192.168.200.1 active Down BGP
Notification received
%BGP_SESSION-5-ADJCHANGE: neighbor 192.168.200.1 IPv4 Unicast
topology base removed from session  BGP Notification received
```

Refer to the exhibit. An engineer noticed that the router log messages do not have any information about when the event occurred. Which action should the engineer take when enabling service time stamps to improve the logging functionality at a granular level?

A. Configure the debug uptime option.

B. Configure the msec option.

C. Configure the timezone option.

D. Configure the log uptime option.

Correct Answer: *B*

Zizu007 Highly Voted 11 months, 2 weeks ago

Selected Answer: B

Correct!

"granular level" in this case ------> msec

upvoted 5 times

GReddy2323 10 months, 3 weeks ago

the granular part went over my head, thank you for the clarification.

upvoted 1 times

inteldarvid Most Recent 5 months, 1 week ago

Selected Answer: B

yes is B:
https://www.cisco.com/E-Learning/bulk/public/tac/cim/cib/using_cisco_ios_software/cmdrefs/service_timestamps.htm

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: B

service timestamps log datetime msec localtime

upvoted 1 times

Question #370

Topic 1

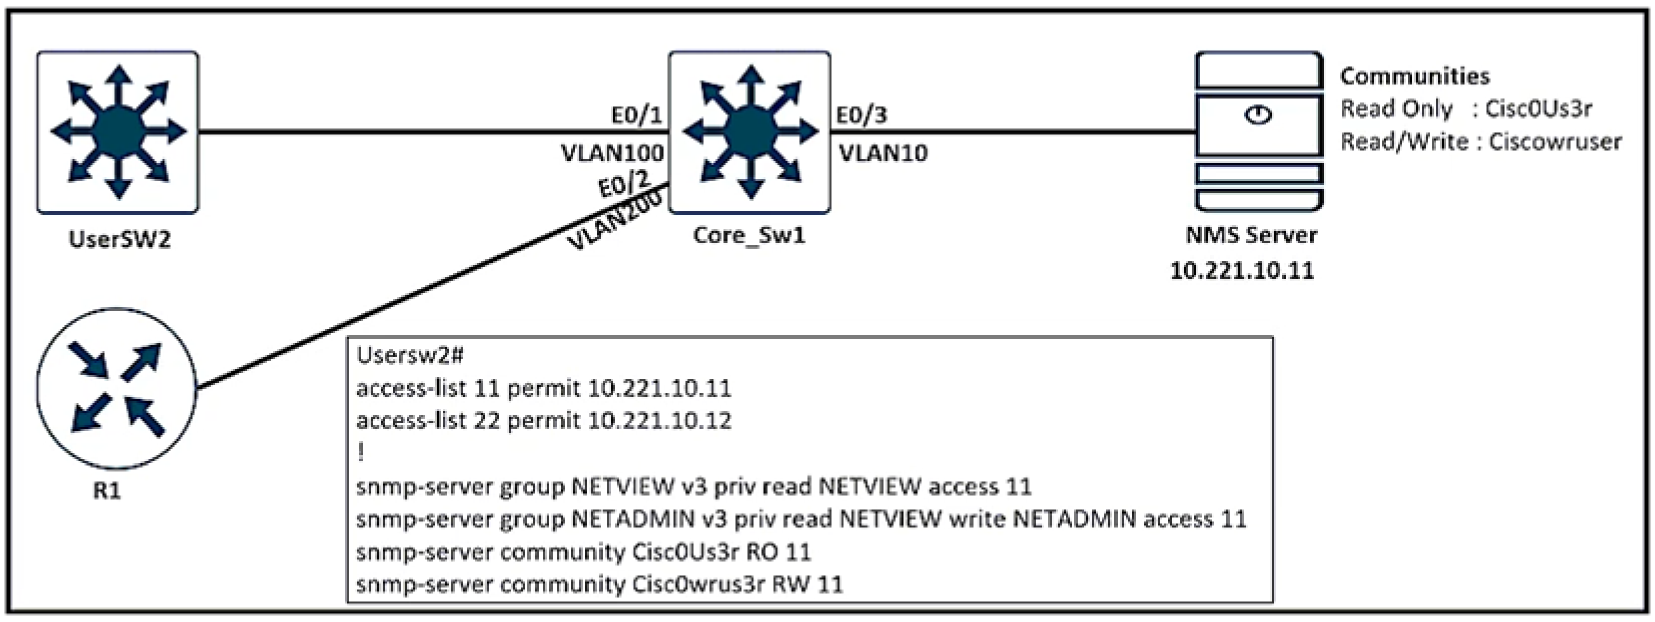

```
Usersw2#
access-list 11 permit 10.221.10.11
access-list 22 permit 10.221.10.12
!
snmp-server group NETVIEW v3 priv read NETVIEW access 11
snmp-server group NETADMIN v3 priv read NETVIEW write NETADMIN access 11
snmp-server community Cisc0Us3r RO 11
snmp-server community Cisc0wrus3r RW 11
```

Refer to the exhibit. An engineer configured SNMP Communities on UserSW2 switch, but the SNMP server cannot upload modified configurations to the switch. Which configuration resolves this issue?

A. snmp-server community CiscoUs3r RW 11

B. snmp-server community Ciscowruser RW 11

C. snmp-server group NETADMIN v3 priv read NETVIEW write NETADMIN access 22

D. snmp-server group NETVIEW v2c priv read NETVIEW access 11

Correct Answer: *B*

inteldarvid 5 months, 1 week ago

Selected Answer: B

B correct

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: B

The string "Ciscowruser" which is a community string (similar to a per-shared key) should be corrected.
Confusing components: the question gives the same name to a server-group and to a view. Plus, the question names the password-like strings to some sort of users.

Examples:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/configuration/xe-3se/3850/snmp-xe-3se-3850-book/nm-snmp-snmpv3.html
https://www.cbtnuggets.com/blog/technology/networking/how-to-configure-snmpv3-and-how-it-works

upvoted 2 times

mitosenoriko 11 months, 2 weeks ago

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/command/nm-snmp-cr-book/nm-snmp-cr-s5.html

upvoted 3 times

Question #371 Topic 1

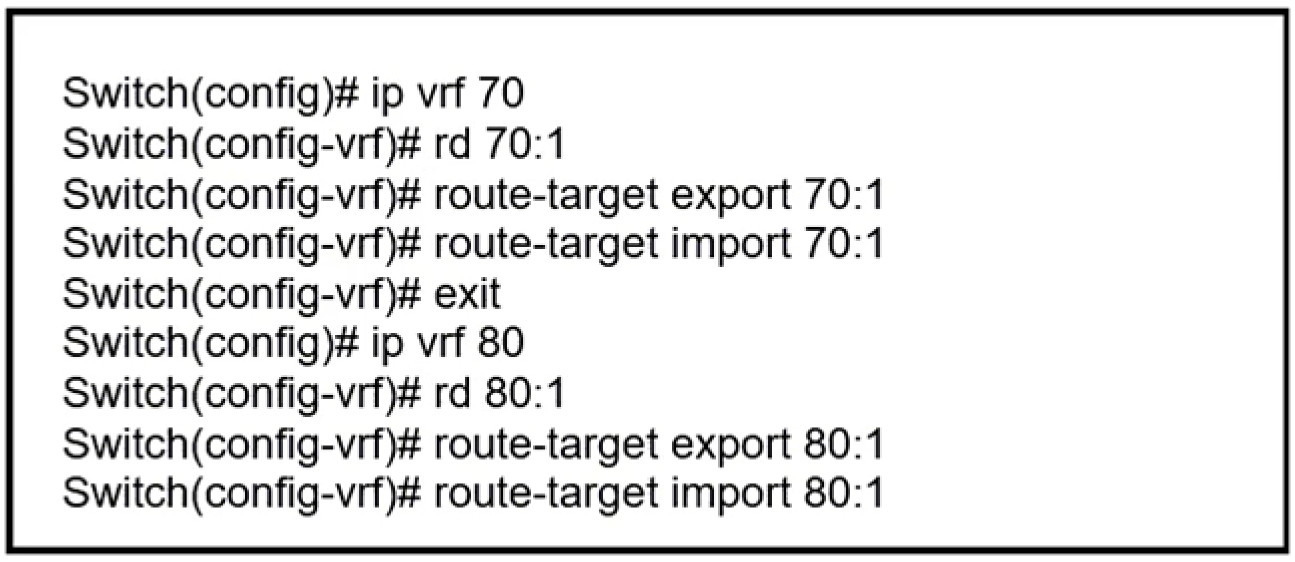

```
Switch(config)# ip vrf 70
Switch(config-vrf)# rd 70:1
Switch(config-vrf)# route-target export 70:1
Switch(config-vrf)# route-target import 70:1
Switch(config-vrf)# exit
Switch(config)# ip vrf 80
Switch(config-vrf)# rd 80:1
Switch(config-vrf)# route-target export 80:1
Switch(config-vrf)# route-target import 80:1
```

Refer to the exhibit. An engineer must extend VRF-Lite over a trunk to another switch for VLAN 70 (10.70.70.0/24) on port GigabitEtheret0/0 and VLAN 80 (10.80.80.0/24) on port GigabitEthernet0/1. Which configuration accomplishes this objective?

A. interface GigabitEthernet0/0
no switchport
ip vrf forwarding 70

ip address 10.70.70.1 255.255.255.0
!
interface GigabitEthernet0/1
no switchport
ip vrf forwarding 80
ip address 10.80.80.1 255.255.255.0

B. interface GigabitEthernet0/0
switchport trunk encapsulation dot1q
switchport mode trunk
switchport trunk allowed vlan 70
!
interface GigabitEthernet0/1
switchport trunk encapsulation dot1q
switchport mode trunk
switchport trunk allowed vlan 80

C. interface GigabitEthernet0/0
switchport mode access
switchport access vlan 70
ip vrf forwarding 70
!
interface GigabitEthernet0/1
switchport mode access
switchport access vlan 80
ip vrf forwarding 80

D. interface GigabitEthernet0/0
switchport mode access
switchport access vlan 70
!
interface GigabitEthernet0/1
switchport mode access
switchport access vlan 80
!

Correct Answer: *A*

Jey117 2 months, 1 week ago

This is stupid. A brakes the trunk from the logical perspective,, you mover it from L2 to a routed port.
B is partially correcto. Because you brake the L2 as You kick out all other vlans.... The command does nota Say "add"

upvoted 1 times

Chiaretta 5 months ago

Selected Answer: A

Given answer is correct. The interface is phisical not sub interface. For this reason the dot1q protocol VLAN trunking is not necessary.

upvoted 2 times

inteldarvid 5 months, 1 week ago

Selected Answer: B

option B correct

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: B

Example configs:
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2/31sg/configuration/guide/conf/vrf.html
https://community.cisco.com/t5/switching/vrf-routing-in-vlan/td-p/3194751

The physical trunk ports are configured with simple trunk configuration. The vrf is associated under the vlan interfaces.
1)create vlans
2)create vrf definition
3)create vlan interface, assign vrf, then IP
4)configure the trunk and allow vlans as required

upvoted 3 times

Typovy 8 months ago

Selected Answer: B

B is correct

upvoted 2 times

6dd4aa0 8 months, 2 weeks ago

Selected Answer: B

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2/20ewa/configuration/guide/conf/vrf.pdf

Look at page 26-7 onwards

upvoted 1 times

[Removed] 9 months, 3 weeks ago

Selected Answer: B

layer 2 is not VRF, it is dot1q. So trunk away like normal and let the L3 device assign the VRF to the l3 interface.

upvoted 4 times

Lilienen 10 months, 2 weeks ago

Could someone explain this question? It clearly says 'over a trunk', but how does one extend VRF over a trunk? Cisco is mixing routing and switching here.
Shouldn't the answer be B, as that's the only answer with an actual trunk configuration?

upvoted 4 times

ellen_AA 11 months, 1 week ago

Selected Answer: A

Given answer is correct

upvoted 2 times

Question #372 *Topic 1*

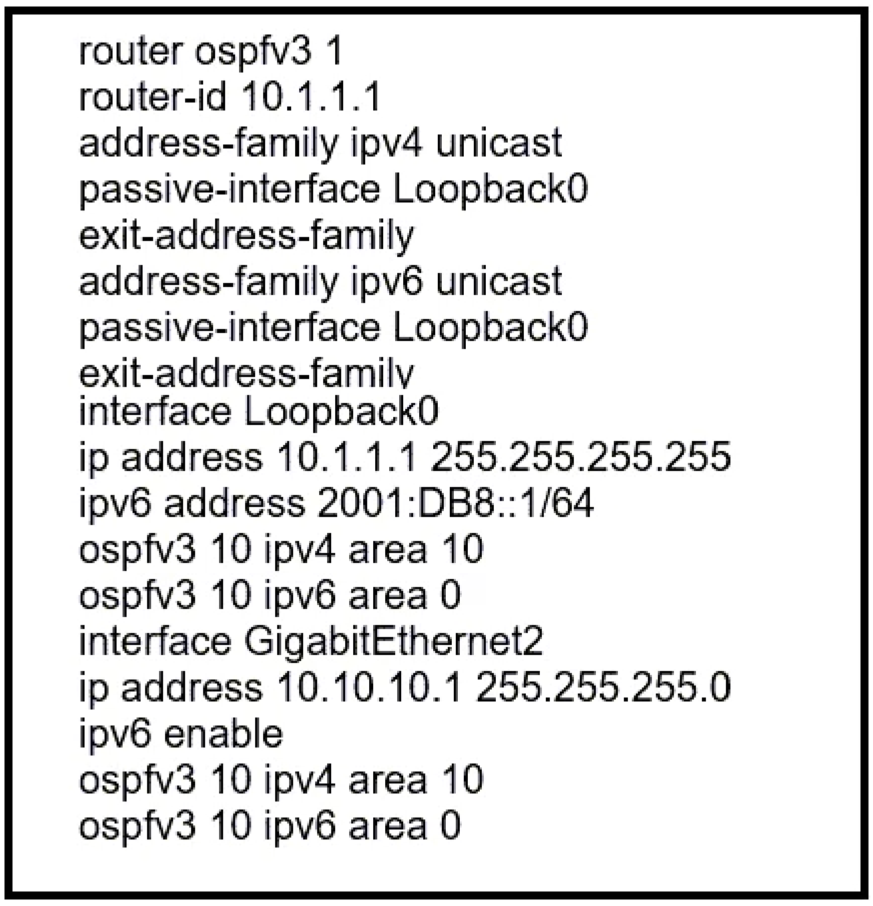

```
router ospfv3 1
router-id 10.1.1.1
address-family ipv4 unicast
passive-interface Loopback0
exit-address-family
address-family ipv6 unicast
passive-interface Loopback0
exit-address-family
interface Loopback0
ip address 10.1.1.1 255.255.255.255
ipv6 address 2001:DB8::1/64
ospfv3 10 ipv4 area 10
ospfv3 10 ipv6 area 0
interface GigabitEthernet2
ip address 10.10.10.1 255.255.255.0
ipv6 enable
ospfv3 10 ipv4 area 10
ospfv3 10 ipv6 area 0
```

Refer to the exhibit. An administrator must configure the router with OSPF for IPv4 and IPv6 networks under a single process. The OSPF adjacencies are not established and did not meet the requirement. Which action resolves the issue?

A. Replace OSPF process 10 on the interfaces with OSPF process 1 for the IPv4 address, and remove process 10 from the global configuration.

B. Replace OSPF process 10 on the interfaces with OSPF process 1, and configure an additional router ID with IPv6 address.

C. Replace OSPF process 10 on the interfaces with OSPF process 1, and remove process 10 from the global configuration.

D. Replace OSPF process 10 on the interfaces with OSPF process 1 for the IPv6 address, and remove process 10 from the global configuration.

Correct Answer: *D*

diegodavid82 4 months ago

Selected Answer: C

is "C" only one process

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

is "C"

upvoted 1 times

DenskyDen 6 months, 3 weeks ago

Selected Answer: C

See HungarianDish comment.

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: C

It should be OSPF process 1 under the interface config for both address families (ipv4, ipv6).

upvoted 4 times

sajjad_gayyem 6 months, 3 weeks ago

Should be c because the question want them under single process

upvoted 1 times

Maholi 7 months, 2 weeks ago

Correct answer is C

upvoted 1 times

forccnp 10 months ago

Why D?
Correct answer is C:/

upvoted 1 times

Zizu007 11 months, 3 weeks ago

Selected Answer: C

as "JKStinn" commented 'C'

upvoted 3 times

JKStinn 1 year ago

Should be C.

upvoted 1 times

Question #373

Topic 1

What is the purpose of an OSPF sham-link?

A. to allow inter-area routing when OSPF is used as the PE-CE connection protocol in an MPLS VPN network

B. to allow intra-area routing when OSPF is used as the PE-CE connection protocol in an MPLS VPN network

C. to correct OSPF backdoor routing when OSPF is used as the PE-CE connection protocol in an MPLS VPN network

D. to correct OSPF backdoor routing when OSPF is used as the PE-PE connection protocol in an MPLS VPN network

Correct Answer: *C*

DUBC89x Highly Voted 1 year ago

Selected Answer: C

OSPF Sham Links are required when you try to use a backdoor link between two CE routers in an MPLS VPN PE CE scenario where you use OSPF as the PE-CE routing protocol. This is best explained with an example, take a look at the following topology:

upvoted 8 times

Titini Highly Voted 10 months, 1 week ago

Selected Answer: D

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/xe-16/iro-xe-16-book/iro-sham-link.html

In an MPLS network, the Provider Edge (PE) routers are typically located in different OSPF areas, with the Provider (P) routers in between them. To enable OSPF routing between the PE routers, a sham-link is created across the MPLS network, connecting the two PE routers and simulating a direct link between them.

upvoted 8 times

inteldarvid 5 months, 1 week ago

You are wrong my friend la awser correct is C:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/xe-16/iro-xe-16-book/iro-sham-link.html

upvoted 1 times

GReddy2323 9 months, 2 weeks ago

In that same document:
A sham-link overcomes the OSPF default behavior for selecting an intra-area backdoor route between VPN sites instead of an interarea (PE-to-PE) route. A sham-link ensures that OSPF client sites that share a backdoor link can communicate over the MPLS VPN backbone and participate in VPN services.

upvoted 1 times

cebra Most Recent 2 weeks, 3 days ago

Selected Answer: D

https://i2.wp.com/ipwithease.com/wp-content/uploads/2017/12/010-understand-configure-ospf-sham-links01.png?resize=700%2C518&ssl=1

upvoted 1 times

Mohammad963 3 months, 1 week ago

Selected Answer: D

The sham link is a logical link, similar to a virtual link. It allows you to create a point-to-point connection between the two PE routers. The PE routers are then able to flood LSAs across the MPLS VPN backbone. You don't have to configure anything on the CE routers.

upvoted 1 times

inteldarvid 5 months ago

Selected Answer: C

team 100% is C
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/xe-16/iro-xe-16-book/iro-sham-link.html#GUID-663C82AA-E6A1-4D95-9EAE-ED1E41E43F78

upvoted 2 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

team look this pelase, all you: (when use routing ospf, is PE-CE)
https://networklessons.com/mpls/mpls-layer-3-vpn-pe-ce-ospf-sham-link

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

C correct:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/xe-16/iro-xe-16-book/iro-sham-link.html

upvoted 1 times

sajjad_gayyem 6 months ago

Selected Answer: C

Answer is c, the shame link is configured between 2 PE routers to correct the backdoor routing between PE-CE.

check below for more description.
https://networklessons.com/mpls/mpls-layer-3-vpn-pe-ce-ospf-sham-link

upvoted 1 times

Juraj22 6 months, 1 week ago

Selected Answer: D

Correct is D, between PE - PE, check pictures on google

upvoted 1 times

sajjad_gayyem 6 months ago

Yes sham link is configured between PE-PE to solve the problem of routing , but the OSPF is configured between PE-CE.

upvoted 2 times

pyrokar 7 months ago

Selected Answer: C

"Although OSPF PE-CE connections assume that the only path between two client sites is across the MPLS VPN backbone, backdoor paths between VPN sites (shown in grey in the figure below) may exist. If these sites belong to the same OSPF area, the path over a backdoor link will always be selected because OSPF prefers intraarea paths to interarea paths. (PE routers advertise OSPF routes learned over the VPN backbone as interarea paths.) For this reason, OSPF backdoor links between VPN sites must be taken into account so that routing is performed based on policy. "

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/xe-16/iro-xe-16-book/iro-sham-link.html

upvoted 2 times

HungarianDish 7 months, 1 week ago

Selected Answer: C

What is the purpose?
C) to correct OSPF backdoor routing when OSPF is used as the PE-CE connection protocol in an MPLS VPN network
(+backdoor link exists between the ospf sites, +it should be used only for backup purposes)
How does it achieve that?
D) OSPF is used as the PE-PE connection protocol

upvoted 2 times

Malasxd 7 months, 2 weeks ago

Selected Answer: C

C and D are exacly the same answer....

upvoted 1 times

Malasxd 7 months ago

Sorry, they are not the same. C is right

upvoted 3 times

azzawim 8 months, 4 weeks ago

Selected Answer: D

Correct answer D

upvoted 1 times

Slippydippy 9 months ago

Selected Answer: C

In an MPLS VPN configuration, the OSPF protocol is one way you can connect customer edge (CE) routers to service provider edge (PE) routers in the VPN backbone. OSPF is often used by customers who run OSPF as their intrasite routing protocol, subscribe to a VPN service, and want to exchange routing information between their sites using OSPF (during migration or on a permanent basis) over an MPLS VPN backbone.

upvoted 1 times

forccnp 9 months ago

Selected Answer: C

It's C

upvoted 2 times

sasasan12345 9 months, 1 week ago

Selected Answer: D

D is correct.

upvoted 1 times

TAZZER 11 months, 1 week ago

I think the answer could be D because this link is between PE-PE or CE and CE, check out the URL
https://ipwithease.com/understand-configure-ospf-sham-links/

upvoted 2 times

Question #374

Topic 1

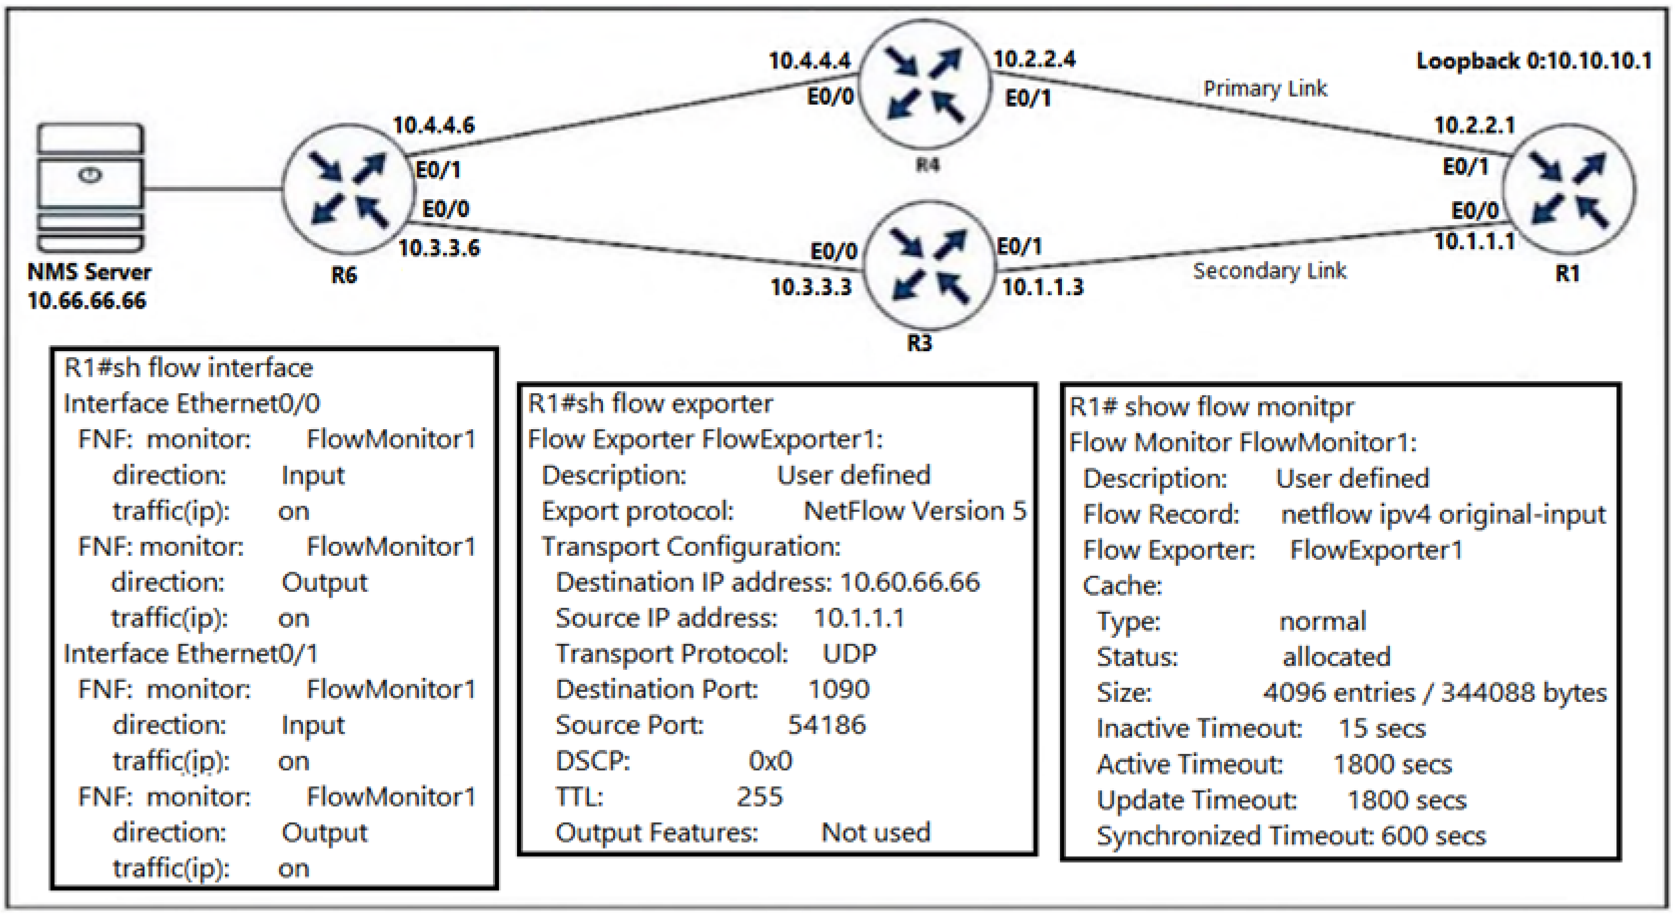

```
R1#sh flow interface
Interface Ethernet0/0
  FNF:  monitor:      FlowMonitor1
        direction:    Input
        traffic(ip):  on
  FNF: monitor:       FlowMonitor1
        direction:    Output
        traffic(ip):  on
Interface Ethernet0/1
  FNF:  monitor:      FlowMonitor1
        direction:    Input
        traffic(ip):  on
  FNF:  monitor:      FlowMonitor1
        direction:    Output
        traffic(ip):  on
```

```
R1#sh flow exporter
Flow Exporter FlowExporter1:
  Description:          User defined
  Export protocol:        NetFlow Version 5
  Transport Configuration:
    Destination IP address: 10.60.66.66
    Source IP address:     10.1.1.1
    Transport Protocol:    UDP
    Destination Port:     1090
    Source Port:         54186
    DSCP:             0x0
    TTL:             255
    Output Features:      Not used
```

```
R1# show flow monitpr
Flow Monitor FlowMonitor1:
  Description:      User defined
  Flow Record:      netflow ipv4 original-input
  Flow Exporter:    FlowExporter1
  Cache:
    Type:            normal
    Status:          allocated
    Size:          4096 entries / 344088 bytes
    Inactive Timeout:    15 secs
    Active Timeout:     1800 secs
    Update Timeout:      1800 secs
    Synchronized Timeout: 600 secs
```

Refer to the exhibit. An engineer configured NetFlow on R1, but the flows do not reach the NMS server from R1. Which configuration resolves this issue?

A. R1(config)#flow monitor FlowMonitort1
R1(config-flow-monitor)#destination 10.66.66.66

B. R1(config)#interface Ethernet0/0
R1(config-if)#ip flow monitor Flowmonitor1 input
R1(config-if)#ip flow monitor Flowmonitor1 output

C. R1(config)#interface Ethernet0/1
R1(config-if)#ip flow monitor Flowmonitor1 input
R1(config-if)#ip flow monitor Flowmonitor1 output

D. R1(config)#flow exporter FlowExporter1
R1(config-flow-exporter)#destination 10.66.66.66

Correct Answer: *D*

forccnp 10 months ago

Selected Answer: D

Given answer is correct

upvoted 2 times

ellen_AA 11 months, 1 week ago

Given answer is correct, netflow destination is configured under the flowexporter

upvoted 1 times

Question #375 Topic 1

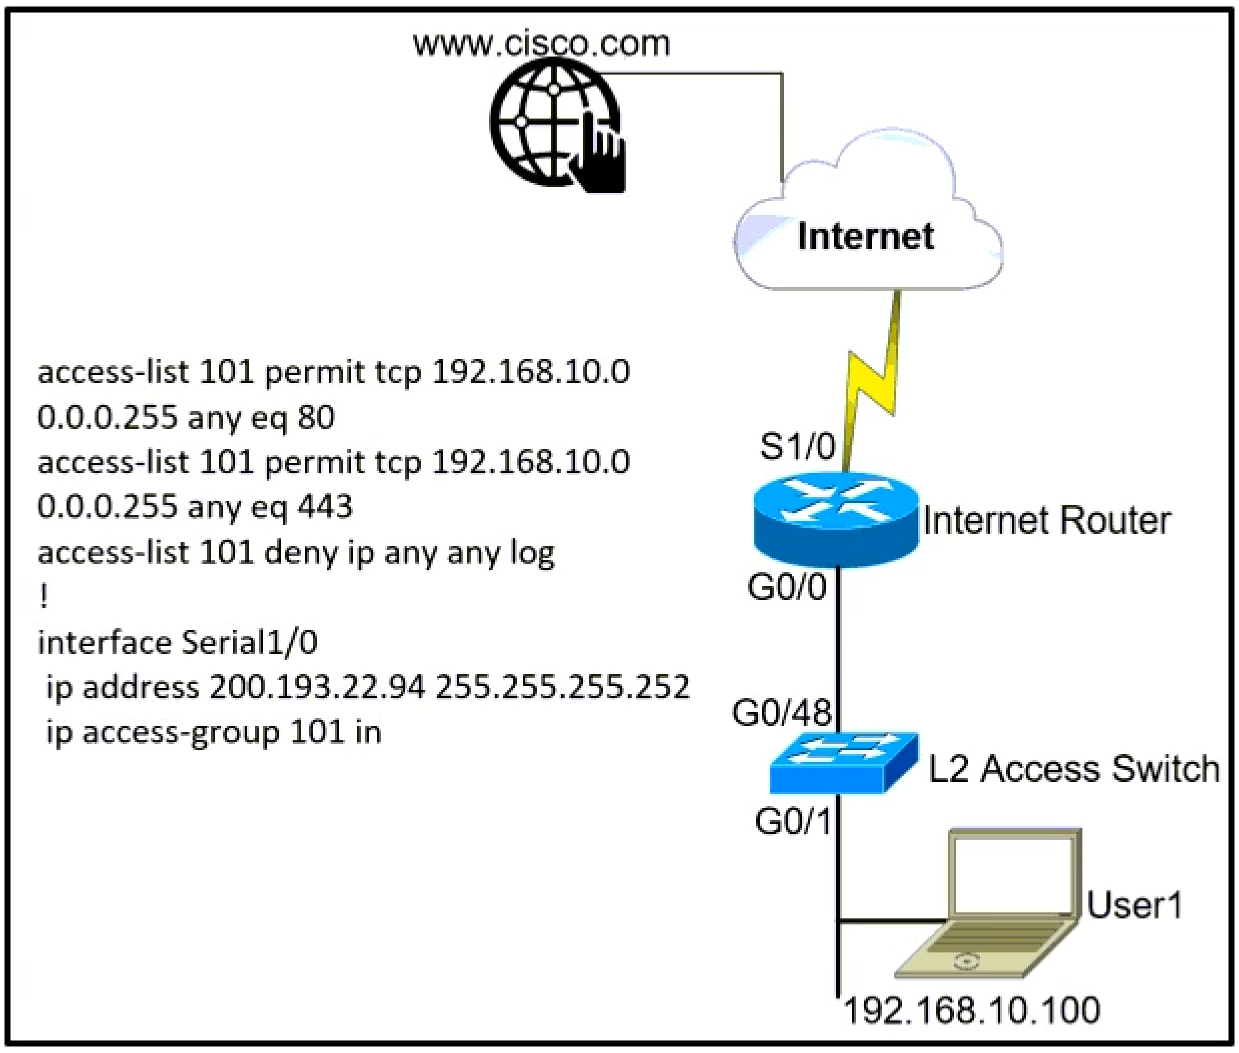

```
access-list 101 permit tcp 192.168.10.0
0.0.0.255 any eq 80
access-list 101 permit tcp 192.168.10.0
0.0.0.255 any eq 443
access-list 101 deny ip any any log
!
interface Serial1/0
 ip address 200.193.22.94 255.255.255.252
 ip access-group 101 in
```

Refer to the exhibit. A network administrator is tasked to permit http and https traffic only toward the internet from the User1 laptop to adhere to company's security policy. The administrator can still ping to www.cisco.com. Which interface should the access list 101 be applied to resolve this issue?

A. Interface G0/0 in the outgoing direction.

B. Interface G0/0 in the incoming direction.

C. Interface S1/0 in the outgoing direction.

D. Interface G0/48 in the incoming direction.

Correct Answer: *B*

b8os5h 1 month ago

Selected Answer: B

NAT should be enabled.In the S1/0 output direction, the ACL is after SNAT conversion (192.168.10.0 -> 200.193.22.94), so it does not match the ACL101 condition.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: B

Correct B

upvoted 1 times

puipuimolcar 7 months, 1 week ago

Why is C incorrect?

upvoted 1 times

RouterToRooter 3 weeks, 1 day ago

That will stop the administrator from being able to ping Cisco.com

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

C is not incorrect, just not that optimal, as traffic that should be blocked by the ACL anyways is transiting the router unnecessarily. Making the router use more resources.

upvoted 1 times

bolbolskanes 12 months ago

The given answer is correct, but the question needs to be corrected: "The administrator still CAN NOT ping to www.cisco.com.

upvoted 2 times

⊟ **forccnp** 10 months ago

tasked to permit http and https traffic only toward the internet.
Not icmp

upvoted 3 times

Question #376 *Topic 1*

An engineer configured routing between multiple OSPF domains and introduced a routing loop that caused network instability. Which action resolves the problem?

A. Set a tag using the redistribute command toward a domain and deny inbound in the other domain by a matching tag.

B. Set a tag using the redistribute command toward a different domain and deny the matching tag when exiting from that domain.

C. Set a tag using the network command in a domain and use the route-map command to deny the matching tag when exiting toward a different domain.

D. Set a tag using the network command in a domain and use the route-map command to deny the matching tag when entering into a different domain.

Correct Answer: *A*

ZamanR 6 days, 4 hours ago

A is the answer

upvoted 1 times

night_wolf_in 1 month, 1 week ago

Selected Answer: B

It is Either A or B. However, for the routing to work between the two domains, the tagged routes need to be advertised first.
So, for A, it says we are Denying the tag in to the new domain. While B says, we are Denying the Tag from the other domain coming back to the origin domain.
To me, this is language question and not technical.

upvoted 1 times

Ghauri777 1 month, 3 weeks ago

Selected Answer: A

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/4170-ospfprocesses.html

This looks like the same scenario.

upvoted 1 times

siyamak 4 months ago

The correct answer is B

upvoted 2 times

guy276465281819372 5 months ago

Selected Answer: D

I dont see where it states in the question that the routes are redistributed, I think it is just an adjacency made with the network command. therefore D seems right.

upvoted 1 times

inteldarvid 4 months, 3 weeks ago

bro, your answer is wrong. try it in a lab

upvoted 2 times

inteldarvid 4 months, 3 weeks ago

```
router ospf 1
log-adjacency-changes
redistribute eigrp 1 metric 10 metric-type 1 subnets tag 10 (correct)
!
R1(config-router)#network ?
A.B.C.D Network number
!
R1(config-router)#network 1.1.1.1 ?
A.B.C.D OSPF wild card bits
!
R1(config-router)#network 1.1.1.1 0.0.0.0 area 0 ?
<cr>
!
R1(config-router)#network 1.1.1.1 0.0.0.0 area 0 (the command does not exist)
```

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: A

the the option corerct is A:

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/4170-ospfprocesses.html

```
router ospf 1
redistribute ospf 2 subnet tag 1
distribute-list 1 route-map filter_domain2 in
!
route-map filter_domain2 deny 10
match tag 2
route-map filter_domain2 permit 20

router ospf 2
redistribute ospf 1 subnet tag 2
distribute-list 1 route-map filter_domain1 in
!
route-map filter_domain1 deny 10
match tag 1
route-map filter_domain1 permit 20
```

upvoted 1 times

Stylar 5 months, 2 weeks ago

Selected Answer: D

Most commends forgot to read the question properly it says "between multiple OSPF domains". Since redistribution is used to inject between different routing protocols we will need to configure a route-map that will match and deny the same incoming route from different OSPF domain.

upvoted 2 times

keesu 6 months ago

Selected Answer: D

This is not redistribution between routing protocols, only between ospf domains, hence D.

upvoted 3 times

slcc99 5 months, 2 weeks ago

OSPF Redistribution Among Different OSPF Processes - Filter Routes Based on Tags

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/4170-ospfprocesses.html#anc16:~:text=be%20very%20difficult.-,Filter%20Routes%20Based%20on%20Tags,-There%20is%20a

upvoted 2 times

HungarianDish 6 months, 2 weeks ago

Selected Answer: A

https://networklessons.com/cisco/ccie-routing-switching/troubleshooting-metric-redistribution
Redistribution Rule: Never advertise prefixes from routing protocol X into Y and then back into X.

Redistribution goes from one protocol (or routing domain) into the other protocol (or routing domain) using the redistribute command. It's easier to picture redistribution between different protocols. Answer "A" means something like this:

```
router ospf 1
redistribute eigrp 100 subnets tag 20

route-map FILTER deny
match tag 20
route-map FILTER permit

router eigrp 100
redistribute ospf 1 metric 100000 1 255 1 1500 route-map FILTER
```

1)redistribute into ospf and set the tag
2)redistribute into eigrp and match the tag (with action deny)

Set a tag so, routes originating from a routing domain will not be redistributed back into that domain.

upvoted 4 times

6dd4aa0 8 months, 2 weeks ago

Selected Answer: A

Redistribute Path (See a similar figure 17-6 on Page 678)
EIGRP 100 R1 ---> R2 OSPF 1 (Redistribute into OSPF from EIGRP)
EIGRP 100 R3 <---- R4 OSPF 1 (The same redistribute packets will be introduce into EIGRP from OSPF)

```
R2
====
router ospf 1
redistribute eigrp 100 subnet route-map FROM-EIGRP-TO-OSPF

route-map FROM-EIGRP-TO-OSPF permit
ip address match prefix-list PREFIX-TAG
set tag 100
```

```
ip prefix-list PREFIX seq 10 permit 10.1.1.0/24



R4
===
router eigrp 100
redistribute ospf metric 1000000 100 255 1 1500 route-map FROM-OSPF-TO-EIGRP

route-map FROM-EIGRP-TO-OSPF deny
tag match 100
route-map FROM-EIGRP-TO-OSPF permit
```

upvoted 2 times

Typovy 9 months ago

Selected Answer: A

A is correct

upvoted 1 times

forccnp 9 months, 1 week ago

Selected Answer: B

B is the correct answer

upvoted 3 times

Titini 10 months ago

Why not D?

upvoted 1 times

Titini 10 months ago

I think that the tag should be blocked when exiting the different domain in orders for the routes to become known to the different domain. So i will go with B.

upvoted 2 times

Question #377 *Topic 1*

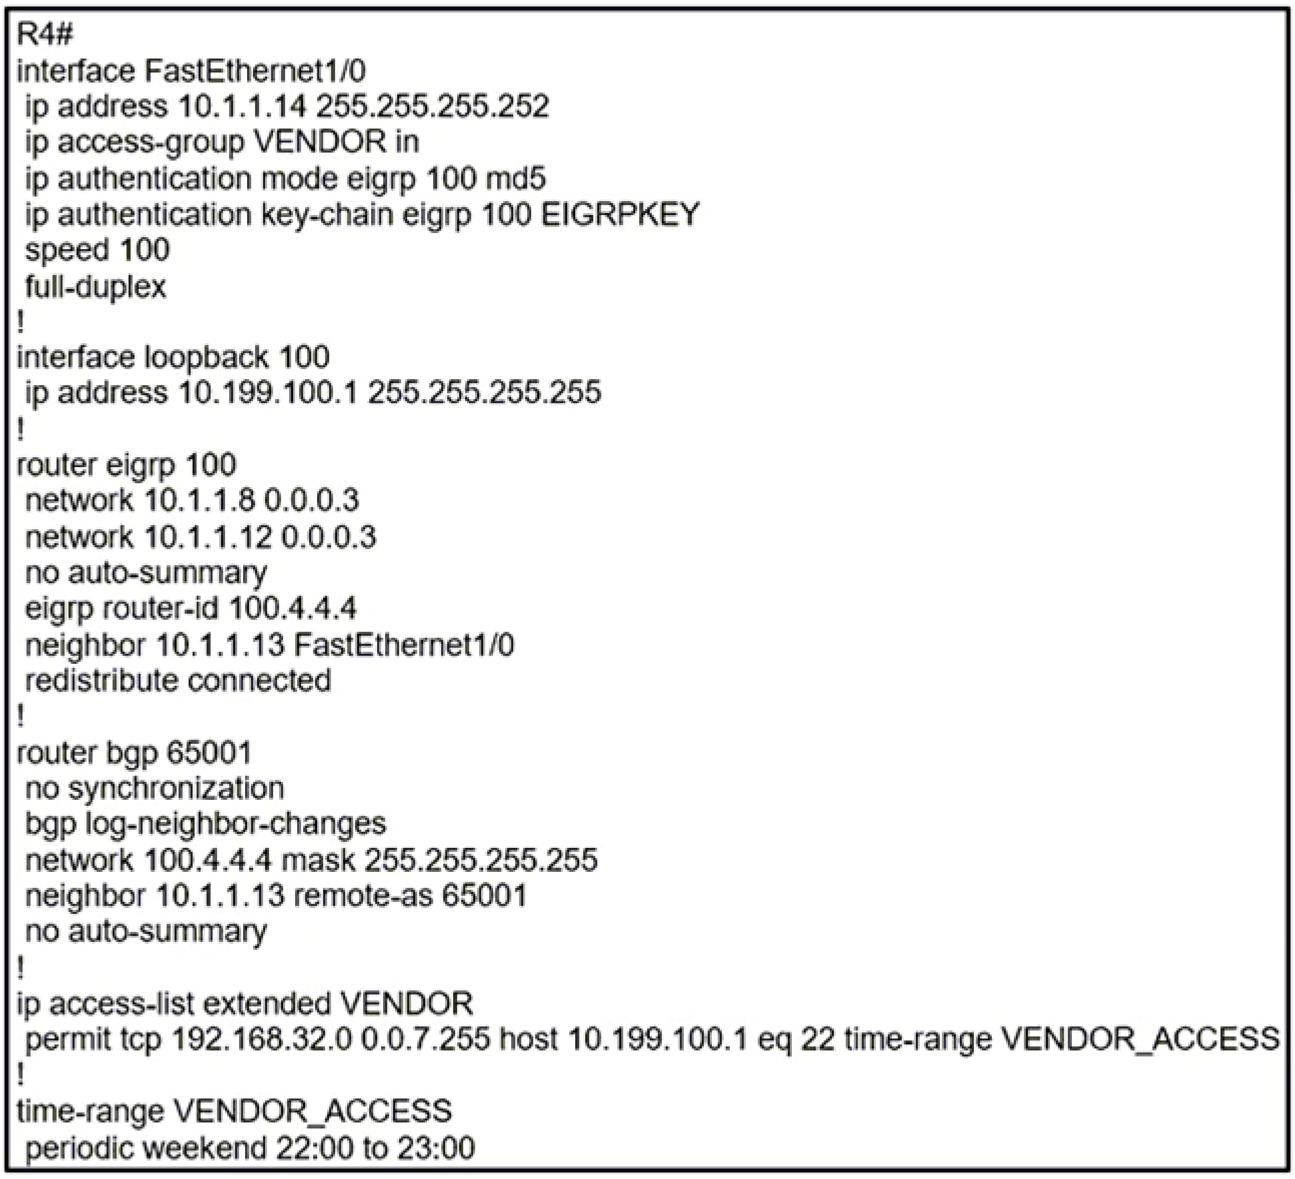

```
R4#
interface FastEthernet1/0
 ip address 10.1.1.14 255.255.255.252
 ip access-group VENDOR in
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 EIGRPKEY
 speed 100
 full-duplex
!
interface loopback 100
 ip address 10.199.100.1 255.255.255.255
!
router eigrp 100
 network 10.1.1.8 0.0.0.3
 network 10.1.1.12 0.0.0.3
 no auto-summary
 eigrp router-id 100.4.4.4
 neighbor 10.1.1.13 FastEthernet1/0
 redistribute connected
!
router bgp 65001
 no synchronization
 bgp log-neighbor-changes
 network 100.4.4.4 mask 255.255.255.255
 neighbor 10.1.1.13 remote-as 65001
 no auto-summary
!
ip access-list extended VENDOR
 permit tcp 192.168.32.0 0.0.7.255 host 10.199.100.1 eq 22 time-range VENDOR_ACCESS
!
time-range VENDOR_ACCESS
 periodic weekend 22:00 to 23:00
```

Refer to the exhibit. A network engineer received a call from the vendor for a failed attempt to remotely log in to their managed router loopback interface from 192.168.40.15. Which action must the network engineer take to resolve the issue?

A. The source IP summarization must be updated to include the vendor source IP address.

B. The time-range configuration must be changed to use absolute instead of periodic.

C. The EIGRP configuration must be updated to include a network statement for loopback 100.

D. The IP access list VENDOR must be applied to interface loopback 100.

Correct Answer: *A*

⊟ **ellen_AA** Highly Voted 11 months, 1 week ago

Selected Answer: A

the extended access-list VENDOR: 192.168.32.0 0.0.7.255 that starts at
192.168.32.1 and ends at 192.168.39.254 doesn't include the vendor ip address of 192.168.40.15

upvoted 6 times

⊟ **AlexInShort12** Most Recent 5 days, 10 hours ago

Selected Answer: C

Are you for real guys?
The loopback is not even advertised. Unless the provider has some static route. Nothing will even go to that loopback.
ACL and summarization are not the same thing.
We are not even sure the ACL is being us to filter the vty access...

upvoted 1 times

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: A

The option correct is A

upvoted 2 times

Question #378 *Topic 1*

DRAG DROP

-

Drag and drop the descriptions from the left onto the corresponding MPLS components on the right.

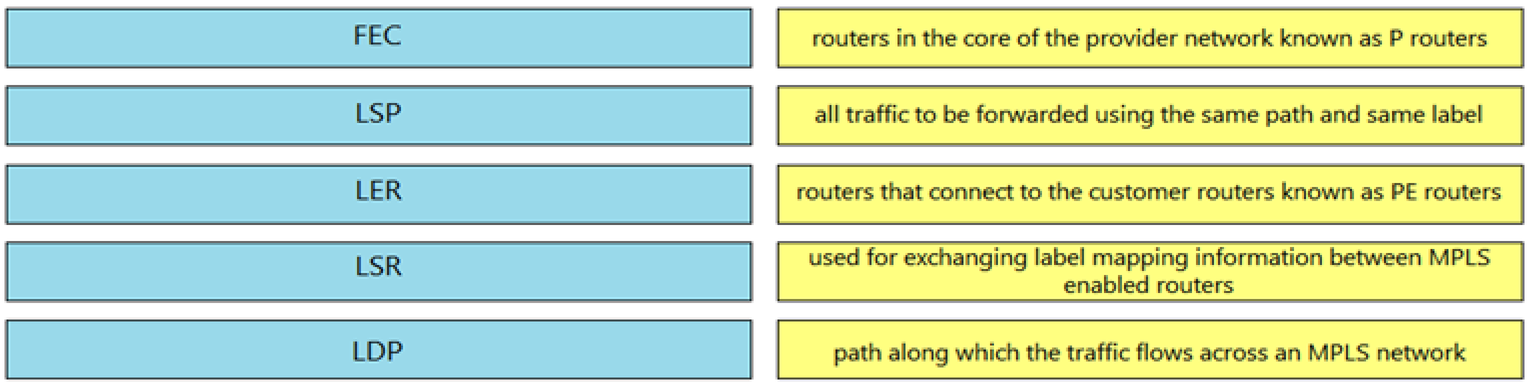

FEC	routers in the core of the provider network known as P routers
LSP	all traffic to be forwarded using the same path and same label
LER	routers that connect to the customer routers known as PE routers
LSR	used for exchanging label mapping information between MPLS enabled routers
LDP	path along which the traffic flows across an MPLS network

Correct Answer:

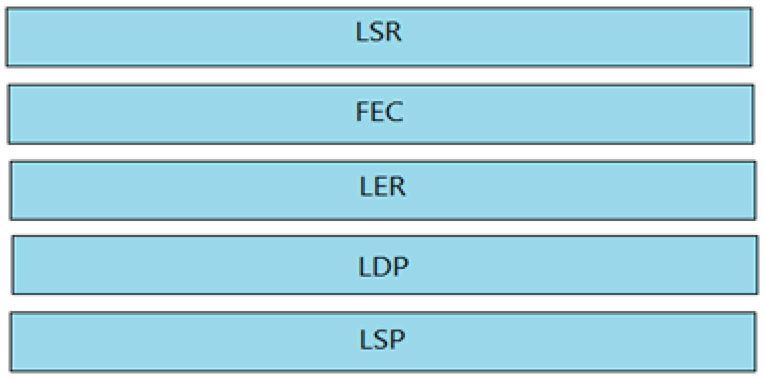

LSR

FEC

LER

LDP

LSP

⊟ **SAMAKEMM** 2 months, 1 week ago

Correct

upvoted 1 times

⊟ **[Removed]** 4 months, 3 weeks ago

Looks good

upvoted 1 times

⊟ **inteldarvid** 5 months, 1 week ago

yes team, correct :)

upvoted 1 times

⊟ **Zizu007** 11 months, 2 weeks ago

Correct!

upvoted 3 times

Question #379 *Topic 1*

Network operations report issues with receiving too many external routes, which caused CPU spike on routers with smaller memories. Which action resolves the issue?

A. Configure the area range command when redistributing on ASBR.

B. Configure the summary-address command when redistributing on ABR.

C. Configure the area range command when redistributing on ABR.

D. Configure the summary-address command when redistributing on ASBR.

Correct Answer: *C*

SAMAKEMM 2 months, 1 week ago

Selected Answer: D

D is correct

upvoted 2 times

inteldarvid 5 months, 1 week ago

Selected Answer: D

yes the option corerct is D

upvoted 1 times

thewhale 6 months, 3 weeks ago

Selected Answer: D

External = ASBR = summary-address

upvoted 4 times

HungarianDish 7 months, 1 week ago

Selected Answer: D

https://www.geeksforgeeks.org/configuring-ospf-route-summarization-in-cisco/
ASBR summary-address <network-id> <prefix-mask> [not-advertise]
ABR area <area-id> range <network-id> <prefix-mask> [advertise |not-advertise]

E1/E2 come from ASBR.

upvoted 1 times

forccnp 9 months, 1 week ago

Selected Answer: D

External routes= ASBR= summary-address command

upvoted 1 times

Zizu007 11 months, 2 weeks ago

Selected Answer: D

"external routes" E1|E2 -----> ASBR (summary-address)

upvoted 1 times

heeeeyajoke 1 year ago

Selected Answer: D

This definitely D, the external and redistributed from the ASBR and the summary-address command is used for summarization. UPDATED

upvoted 2 times

heeeeyajoke 1 year ago

This definitely B, the external and redistributed from the ASBR and the summary-address command is used for summarization

upvoted 1 times

DUBC89x 1 year ago

Selected Answer: D

Based of this information the external routes should be summarized by the ASBR's

upvoted 4 times

DUBC89x 1 year ago

The primary difference between area range and summary-address is in where the command should be applied. Area range should be applied on the ABRs when you are trying to summarize routes between OSPF areas. So area range is used to summarize Type 3 LSAs. On the other hand summary-address should be applied on the ASBRs when you are trying to summarize externally redistributed routes from another protocol domain (eigrp, bgp etc). So summary-address is used to summarize Type5/7 LSAs. One exception to this rule is when you have a NSSA area, the router that is responsible for the conversion of Type7 to Type 5 LSA can have the summary-address applied.

upvoted 3 times

Question #380 Topic 1

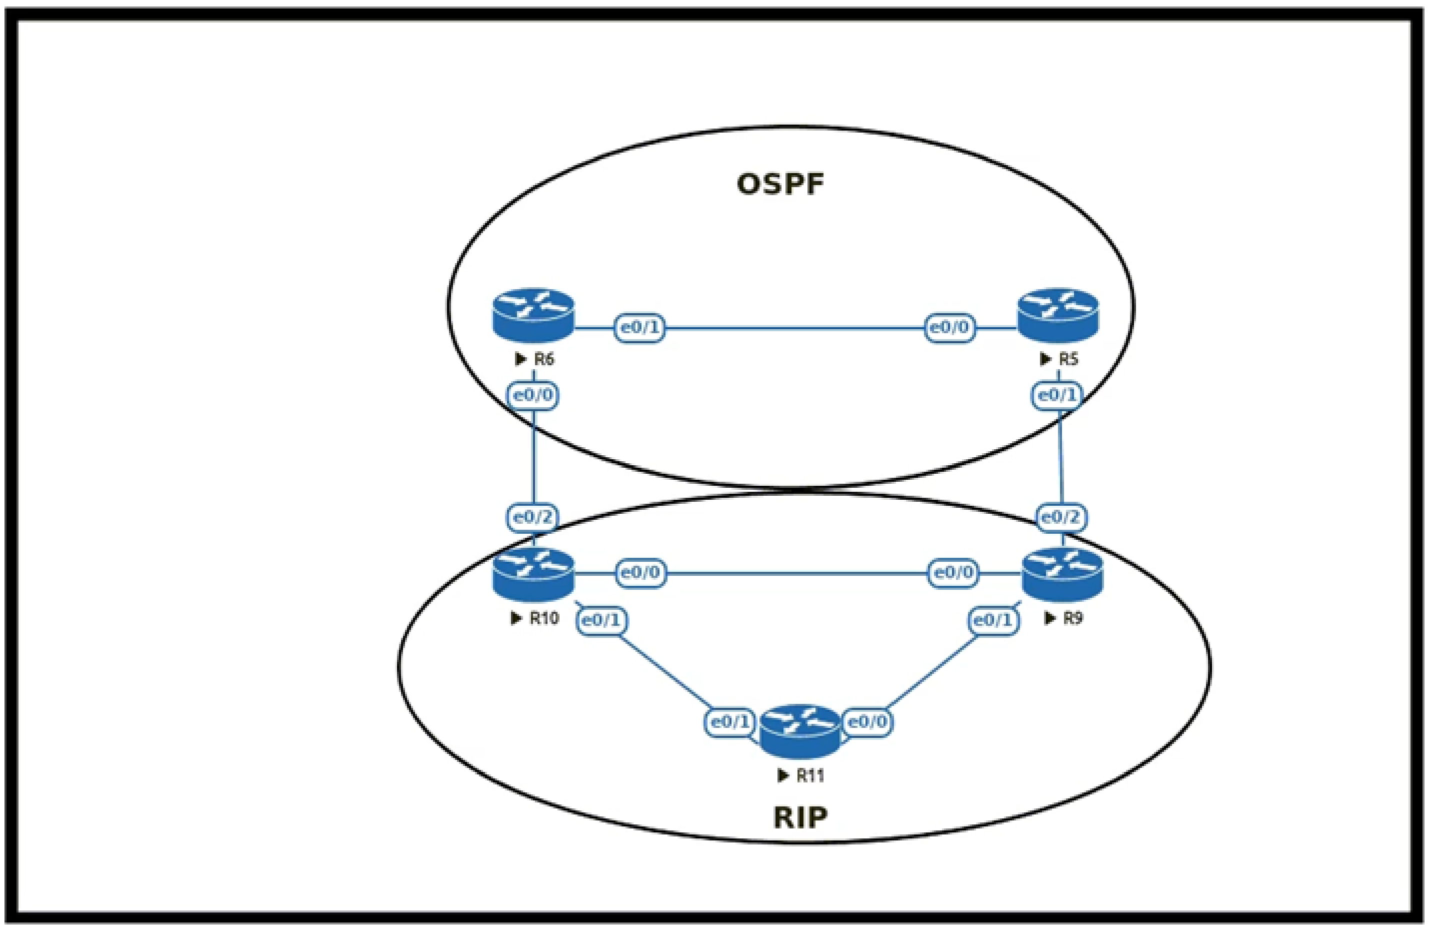

Refer to the exhibit. An engineer must configure OSPF with R9 and R10 and configure redistribution between OSPF and RIP, causing a routing loop. Which configuration on R9 and R10 meets this objective?

```
A. router ospf 1
redistribute rip subnets tag 20
!
route-map deny_tag20 deny 10
match tag 20
route-map deny_tag20 permit 20
!
router ospf 1
distribute-list route-map deny_tag20 in
```

```
B. router ospf 1
redistribute rip subnets tag 20
!
route-map deny_tag20 deny 10
match tag 20
route-map deny_tag20 deny 20
!
router ospf 1
distribute-list route-map deny_tag20 in
```

```
C. router ospf 1
redistribute rip subnets tag 20
!
route-map deny_tag20 deny 10
match tag 20
route-map deny_tag20 permit 20
!
router rip 1
distribute-list route-map deny _tag20 in
```

```
D. router ospf 1
redistribute rip subnets tag 20
!
route-map deny_tag20 permit 10
match tag 20
```

```
route-map deny_tag20 permit 20
!
router ospf 1
distribute-list route-map deny_tag20 in
```

Correct Answer: *C*

fizzer 3 months, 1 week ago

Funny how the omission of a single word can completely change the objective of a sentence
Question:
Refer to the exhibit. An engineer must configure OSPF with R9 and R10 and configure redistribution between OSPF and RIP "without" causing a routing loop. Which configuration on R9 and R10 meets this objective?

Have they erroneously omitted the word "without"? - right answer is A
Have they intentionally omitted the word "without"? - right answer is D

I am just going to chalk this one up to another one of their stupid questions, sometimes they make them because of errors and sometimes they make them trying to intentionally mislead

upvoted 2 times

chris110 3 months, 1 week ago

Selected Answer: D

D.

```
router ospf 1
redistribute rip subnets tag 20
!
route-map deny_tag20 permit 10
match tag 20
route-map deny_tag20 permit 20
!
router ospf 1
distribute-list route-map deny_tag20 in
```

upvoted 1 times

yellowswan 3 months, 2 weeks ago

the description is: "someone MUST configure redistribution and causing a routing loop"
the question is: "which configuration meet this objective"
you are going to select the answer to CAUSE a loop, rather than AVOID a loop

upvoted 1 times

Muste 4 months, 1 week ago

Selected Answer: D

the question says which configuration meets the requirements of causing loop so answer is D

upvoted 1 times

chaocheng 4 months, 1 week ago

```
router ospf 1
redistribute rip subnets tag 20 //set tag as 20
route-map deny_tag20 deny 10 //tag 20，deny this route
match tag 20
route-map deny_tag20 permit 20 // permit other route
router ospf 1
distribute-list route-map deny_tag20 in //use distribute-list filter
```

upvoted 2 times

inteldarvid 5 months, 1 week ago

Sobre, i said , if I avoid a loop is the opción A, and if I create a loop the opción is D. For this quesitos is D. Because the question is explicit create a loop for wrong the engineer

upvoted 2 times

inteldarvid 5 months, 1 week ago

Now if I want to avoid a loop the correct option is A and if I want to create a loop the option is C

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: A

team aswer correct is A:

It has to be in the OSPF domain when we deny the prefix to block upstream in the OSPF domain and avoid the loop (remember that ospf has an AD of 110 and RIP of 120) and the routers prefer a lower AD for OSPF even though they have the downstream network. and a loop is generated. Deny traffic on the destination router in OSPF. because the route is received by external OSPF with a distance AD 110

upvoted 1 times

daloslav 6 months, 4 weeks ago

Selected Answer: A

There is no doubt about that right answer is A - just labbed it.
When you are redistributing from RIP into OSPF, it is redistribution from higher AD (RIP=120) to lower AD (OSPF=110).

R10 receives routes from R11 (RIP routes) -> redistributes it into OPSF to R6 -> R5 -> R9.
At the beginning, R9 see routes from R11 too (RIP routes). When R9 receives that redistributed routes via OSPF, populates its routing table with this routes because OSPF AD (lower than RIP AD) -> this also can happen vice versa on R10 = routing loop.

To break routing loop, you have to filter routes redistributed from RIP to OSPF -> tagg them and do filtering based on tag on both R10 and R9 routers.

upvoted 2 times

HungarianDish 6 months, 2 weeks ago

The distribute-list in "A" causes the redistributed rip routes (tagged 20) blocked from getting into RIB from ospf LSDB table on the specific device where the distribute list is applied. For example, let's see what happens if we apply it on R9:

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

R9 = ASBR. RIP+OSPF mutual redistribution. Follow the rip subnet 11.11.11.11/32.
Before adding the distribute-list:

```
router ospf 1
redistribute rip subnets tag 20
r9(config-router)#
```

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

```
r9(config-router)#do sh ip route | b Gateway
Gateway of last resort is not set

1.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
O E2 1.1.11.0/24 [110/20] via 1.1.59.5, 00:00:31, GigabitEthernet0/0
C 1.1.19.0/24 is directly connected, GigabitEthernet0/3
L 1.1.19.9/32 is directly connected, GigabitEthernet0/3
C 1.1.59.0/24 is directly connected, GigabitEthernet0/0
L 1.1.59.9/32 is directly connected, GigabitEthernet0/0
O 1.1.61.0/24 [110/3] via 1.1.59.5, 00:00:31, GigabitEthernet0/0
O 1.1.65.0/24 [110/2] via 1.1.59.5, 00:00:31, GigabitEthernet0/0
C 1.1.91.0/24 is directly connected, GigabitEthernet0/2
L 1.1.91.9/32 is directly connected, GigabitEthernet0/2
11.0.0.0/32 is subnetted, 1 subnets
O E2 11.11.11.11 [110/20] via 1.1.59.5, 00:00:31, GigabitEthernet0/0
```

=> Injected rip routes by redistribution appear as O E2 routes in routing table.

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

After adding the distribute-list:

```
router ospf 1
redistribute rip subnets tag 20
distribute-list route-map DENY in
r9(config-router)#
```

upvoted 2 times

HungarianDish 6 months, 2 weeks ago

```
r9(config-router)#do sh ip route | b Gateway
Gateway of last resort is not set

1.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
R 1.1.11.0/24 [120/1] via 1.1.91.11, 00:00:04, GigabitEthernet0/2
[120/1] via 1.1.19.10, 00:00:10, GigabitEthernet0/3
C 1.1.19.0/24 is directly connected, GigabitEthernet0/3
L 1.1.19.9/32 is directly connected, GigabitEthernet0/3
C 1.1.59.0/24 is directly connected, GigabitEthernet0/0
L 1.1.59.9/32 is directly connected, GigabitEthernet0/0
O 1.1.61.0/24 [110/3] via 1.1.59.5, 00:00:14, GigabitEthernet0/0
O 1.1.65.0/24 [110/2] via 1.1.59.5, 00:00:14, GigabitEthernet0/0
C 1.1.91.0/24 is directly connected, GigabitEthernet0/2
L 1.1.91.9/32 is directly connected, GigabitEthernet0/2
11.0.0.0/32 is subnetted, 1 subnets
R 11.11.11.11 [120/1] via 1.1.91.11, 00:00:04, GigabitEthernet0/2
r9(config-router)#
```

upvoted 2 times

HungarianDish 6 months, 2 weeks ago

The route 11.11.11.11 (from below example) is still there in ospf LSDB, and gets advertised to the other ospf routers as lsa type 5.
See more info on the topic:
https://networklessons.com/ospf/ospf-distribute-list-filtering

```
r9#show ip ospf database | b External
Type-5 AS External Link States

Link ID ADV Router Age Seq# Checksum Tag
1.1.11.0 1.1.61.10 1463 0x80000002 0x00C272 20
1.1.19.0 1.1.61.10 1463 0x80000002 0x006AC2 20
1.1.19.0 1.1.91.9 1442 0x80000002 0x009D72 20
1.1.91.0 1.1.91.9 1442 0x80000002 0x008245 20
11.11.11.11 1.1.91.9 983 0x80000001 0x008E6B 20
r9#
```

upvoted 1 times

ioreskovic 6 months, 1 week ago

So what if route to 11.11.11.11 remains in LSDB? The goal is to keep RIP 11.11.11.11 route on R9 and R10, isn't it?

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Protocol + route-map combination is incorrect in answer A.

upvoted 1 times

GReddy2323 7 months, 1 week ago

Selected Answer: C

Poorly written question as always, if you read the question it says the engineer must configure redistribution but it caused a routing loop, I don't think they are wanting us to create a routing loop because they are always teaching how to prevent it.

upvoted 3 times

JoeyT 6 months, 1 week ago

objective is to create a loop... like you are testing an interviewer...

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

I agree, probably the question is wrong. Typovy already pointed out that the correct answer would be "C" to avoid the loop. Surprisingly, the question instructs to create a loop, which can be achieved by answer "D".

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

When I tested the configurations in CML, I had a problem with solution "C". Unlike ospf, rip does not have the parameter route-map when using distribution- lists.

```
R10(config)#router rip
R10(config-router)#distribute-list ?
<1-199> IP access list number
<1300-2699> IP expanded access list number
WORD Access-list name
gateway Filtering incoming updates based on gateway
prefix Filter prefixes in routing updates
```

Instead, I needed to choose a classic configuration with redistribution and route-map to match and deny the tagged routes (rip routes coming back from ospf):

```
R10(config)#router rip
R10(config-router)#redistribute ospf 1 route-map MATCH metric 3
```

upvoted 1 times

Typovy 8 months, 1 week ago

Selected Answer: D

The objective is to create routing loop, not avoid loop so answer is D.
C will avoid loop :)

upvoted 4 times

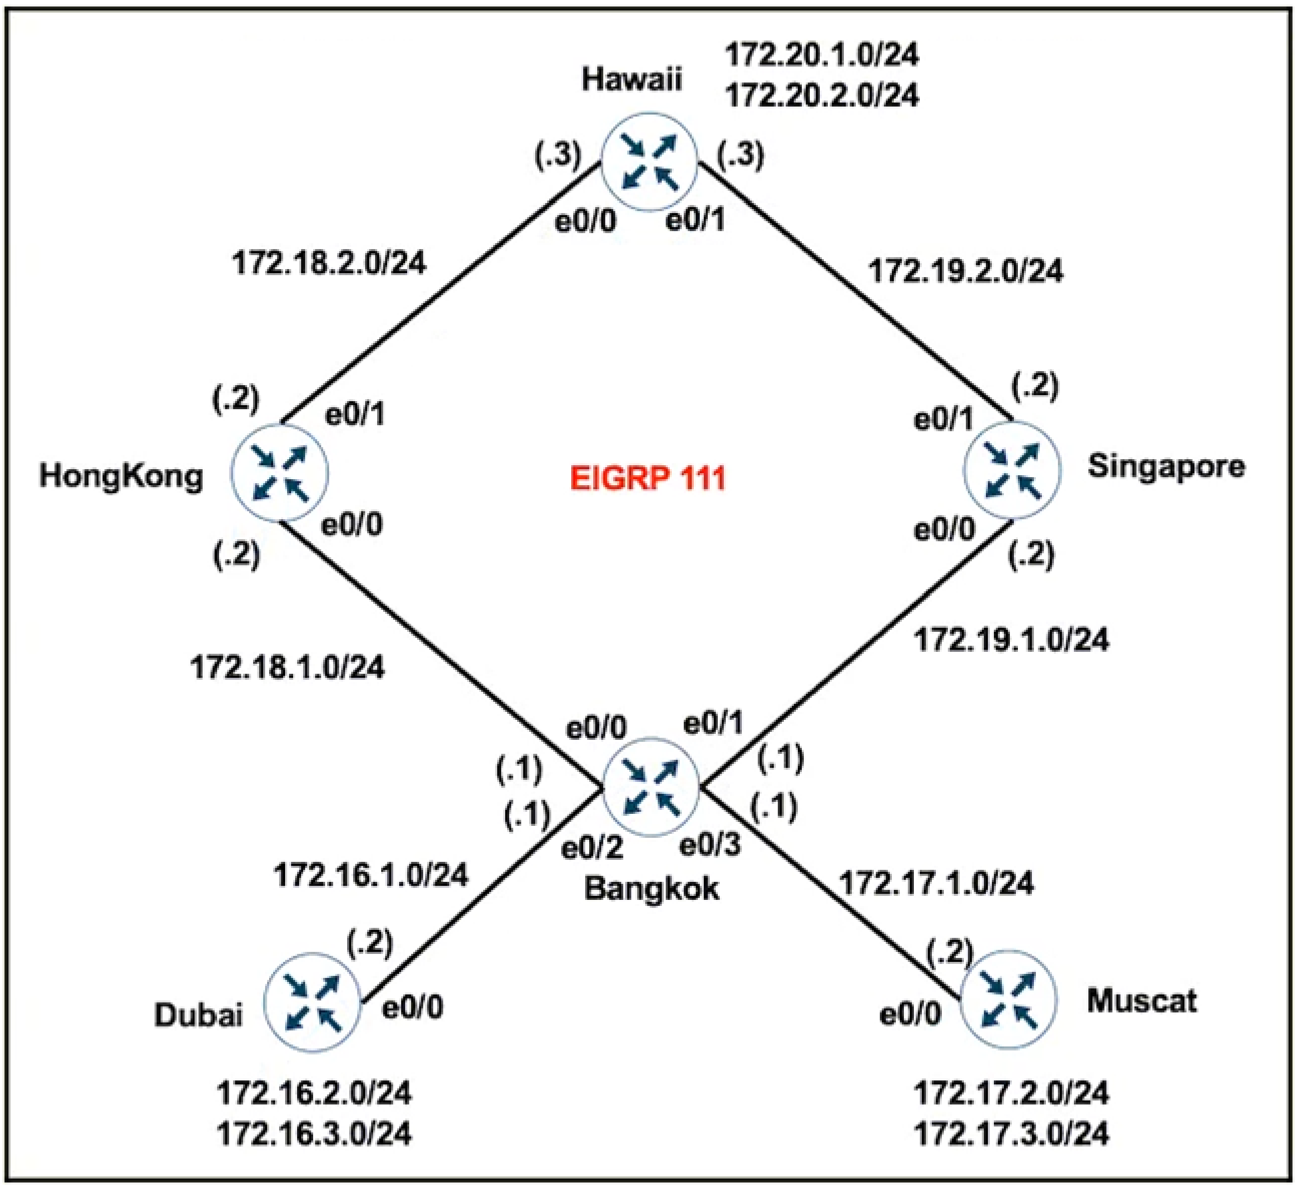

Refer to the exhibit. Bangkok is using ECMP to reach the 172.20.2.0/24 network. The network administrator must configure it in such a way that traffic from 172.16.2.0/24 network uses the Singapore router as the preferred route. Which set of configurations accomplishes this task?

A. Bangkok -

```
access-list 101 permit ip 172.16.2.0 0.0.0.255 172.20.2.0 0.0.0.255
!
route-map PBR1 permit 10
match ip address 101
set ip next-hop 172.19.1.2
!
interface Ethernet0/1
ip policy route-map PBR1
```

B. Dubai -

```
access-list 101 permit ip 172.16.2.0 0.0.0.255 172.20.2.0 0.0.0.255
!
route-map PBR1 permit 10
match ip address 101
set ip next-hop 172.19.1.2
set ip next-hop peer-address
!
interface Ethernet0/0
ip policy route-map PBR1
```

C. Bangkok -

```
access-list 101 permit ip 172.16.2.0 0.0.0.255 172.20.2.0 0.0.0.255
!
route-map PBR1 permit 10
match ip address 101
```

```
set ip next-hop 172.19.1.2
!
interface Ethernet0/2
ip policy route-map PBR1
```

D. Dubai -

```
access-list 101 permit ip 172.16.2.0 0.0.0.255 172.20.2.0 0.0.0.255
!
route-map PBR1 permit 10
match ip address 101
set ip next-hop 172.19.1.2
!
interface Ethernet0/0
ip policy route-map PBR1
```

Correct Answer: *C*

ellen_AA Highly Voted 11 months, 1 week ago

Given answer is correct, PBRs are always applied on the ingress interface

upvoted 5 times

HarwinderSekhon Most Recent 4 months, 1 week ago

Selected Answer: C

PBR on source always preferred

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

yes, its correct, ist similar with AC extenda, always close to the source

upvoted 2 times

Question #382 *Topic 1*

Which label operations are performed by a label edge router?

A. SWAP and POP

B. PUSH and POP

C. SWAP and PUSH

D. PUSH and PHP

Correct Answer: *B*

KungFuPanda19 1 week, 3 days ago

D. Not sure, though; all semantics here. (ENARSI book p739)
Edge router "advertises a pop"; second last router "pops the label".

upvoted 1 times

KungFuPanda19 1 week, 3 days ago

C. Not sure, though; all semantics here. (ENARSI book p739)
Edge router "advertises a pop"; second last router "pops the label".

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: B

yes correct B

https://ipcisco.com/lesson/mpls-label-switching/

upvoted 2 times

Titini 10 months, 1 week ago

Selected Answer: B

B. PUSH and POP: A LER pushes a new label onto incoming packets and pops off the label from outgoing packets. The incoming label is used to determine the forwarding path for the packet, while the outgoing label is used to direct the packet to the next hop in the path.

upvoted 2 times

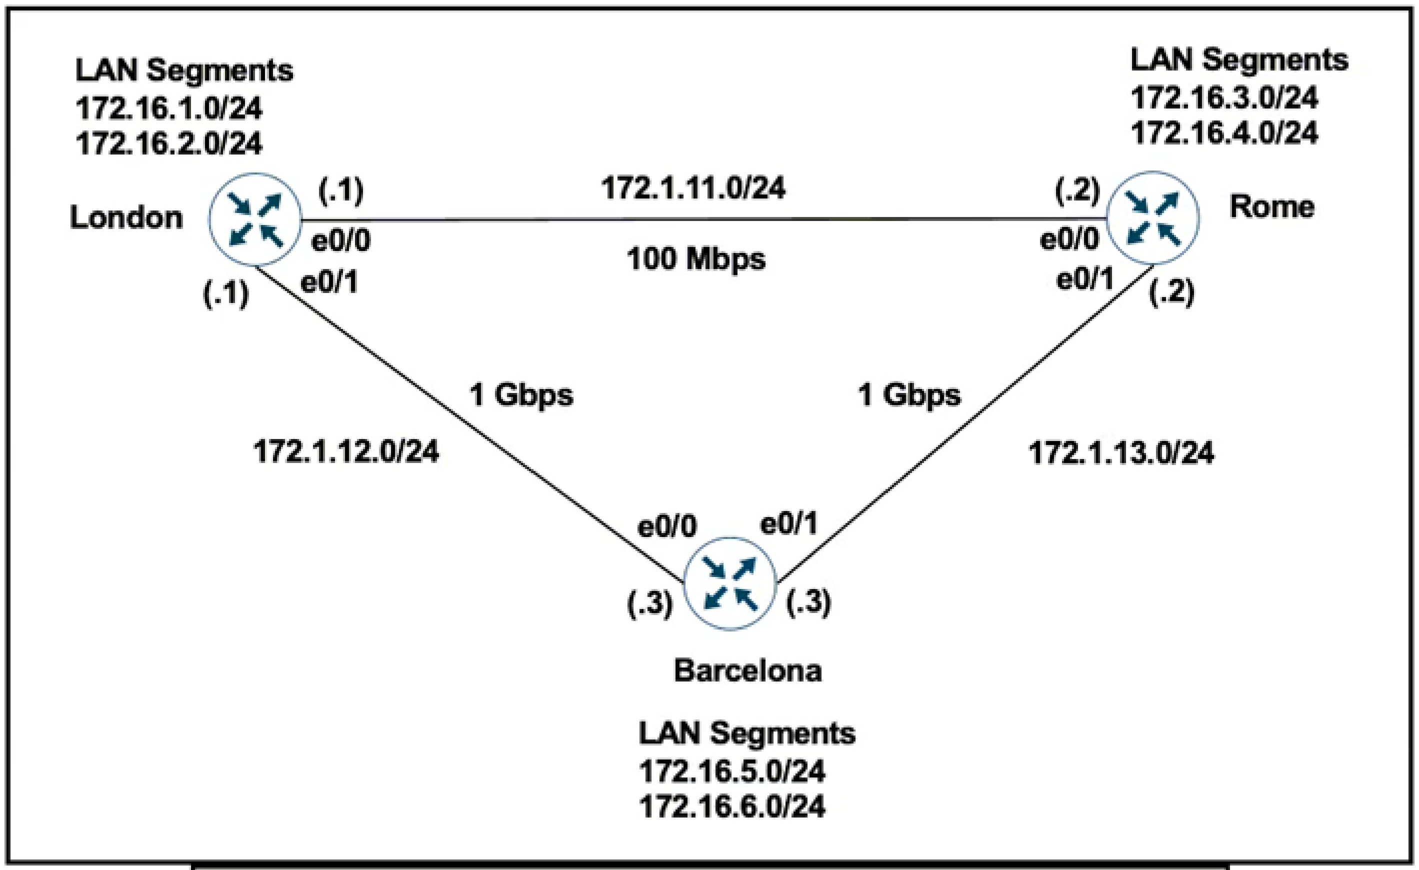

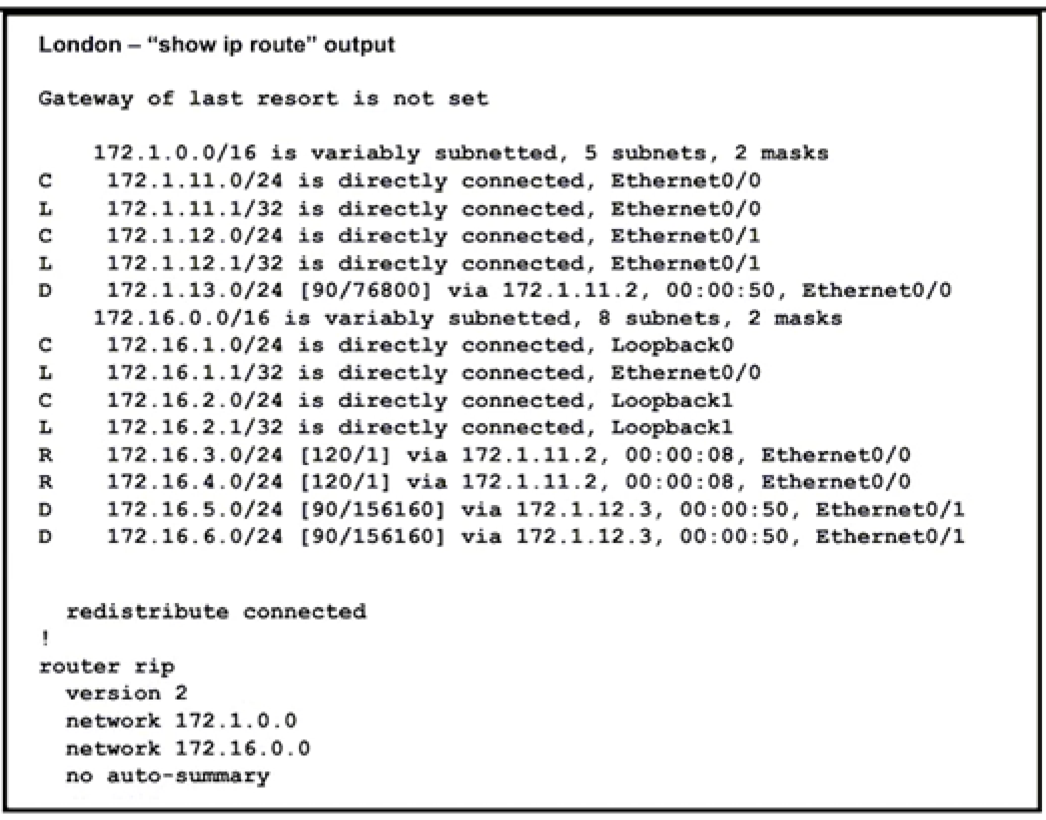

London – "show ip route" output

```
Gateway of last resort is not set

    172.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
C     172.1.11.0/24 is directly connected, Ethernet0/0
L     172.1.11.1/32 is directly connected, Ethernet0/0
C     172.1.12.0/24 is directly connected, Ethernet0/1
L     172.1.12.1/32 is directly connected, Ethernet0/1
D     172.1.13.0/24 [90/76800] via 172.1.11.2, 00:00:50, Ethernet0/0
    172.16.0.0/16 is variably subnetted, 8 subnets, 2 masks
C     172.16.1.0/24 is directly connected, Loopback0
L     172.16.1.1/32 is directly connected, Ethernet0/0
C     172.16.2.0/24 is directly connected, Loopback1
L     172.16.2.1/32 is directly connected, Loopback1
R     172.16.3.0/24 [120/1] via 172.1.11.2, 00:00:08, Ethernet0/0
R     172.16.4.0/24 [120/1] via 172.1.11.2, 00:00:08, Ethernet0/0
D     172.16.5.0/24 [90/156160] via 172.1.12.3, 00:00:50, Ethernet0/1
D     172.16.6.0/24 [90/156160] via 172.1.12.3, 00:00:50, Ethernet0/1


  redistribute connected
!
router rip
  version 2
  network 172.1.0.0
  network 172.16.0.0
  no auto-summary
```

Refer to the exhibits. London must reach Rome using a faster path via EIGRP if all the links are up, but it failed to take this path. Which action resolves the issue?

A. Change the administrative distance of RIP to 150.

B. Increase the bandwidth of the link between London and Barcelona.

C. Use the network statement on London to inject the 172.16.X.0/24 networks into EIGRP.

D. Use the network statement on Rome to inject the 172.16.X.0/24 networks into EIGRP.

Correct Answer: *D*

HungarianDish 7 months, 1 week ago

Selected Answer: D

D is correct, also tested it. If LAN segments (172.16.3.0/24, 172.16.4.0/24) on Rome are advertised by EIGRP then London is choosing the path with the higher bandwidth (minimum BW is 1GB) to reach the LANs on Rome. (Of course, to reach the router Rome, it is choosing the directly connected link with AD 0.)

upvoted 3 times

Zizu007 11 months, 3 weeks ago

Selected Answer: D

correct!
Prefixes from Rome are only advertised in RIP (AD-120). After advertising it in EIGRP these prefixes will be preferred by London site. Additionally, London site performs redistribution of connected routes in EIGRP (see output).

upvoted 3 times

Question #384 *Topic 1*

The network administrator configured the router for Control Plane Policing so that inbound SSH traffic is policed to 500 kbps. This policy must apply to traffic coming in from 10.10.10.0/24 and 192.168.10.0/24 networks.

```
access-list 100 permit ip 10.10.10.0 0.0.0.255 any
access-list 100 permit tcp 192.168.10.0 0.0.0.255 any eq 23
!
class-map CLASS-SSH
match access-group 100
!
policy-map PM-COPP
class CLASS-SSH
police 500000 conform-action transmit
!
interface E0/0
service-policy input PM-COPP
!
interface E0/1
service-policy input PM-COPP
```

The Control Plane Policing is not applied to SSH traffic and SSH is open to use any bandwidth available. Which configuration resolves this issue?

```
A. no access-list 100
access-list 100 permit tcp 10.10.10.0 0.0.0.255 any eq 22
access-list 100 permit tcp 192.168.10.0 0.0.0.255 any eq 22
```

```
B. interface E0/0
no service-policy input PM-COPP
!
interface E0/1
no service-policy input PM-COPP
!
control-plane
service-policy input PM-COPP
```

```
C. no access-list 100
access-list 100 permit tcp 10.10.10.0 0.0.0.255 any eq 22
access-list 100 permit tcp 192.168.10.0 0.0.0.255 any eq 22
!
policy-map PM-COPP
class CLASS-SSH
no police 500000 conform-action transmit
police 500000 conform-action transmit exceed-action drop
```

```
D. no access-list 100
access-list 100 permit tcp 10.10.10.0 0.0.0.255 any eq 22
access-list 100 permit tcp 192.168.10.0 0.0.0.255 any eq 22
!
interface E0/0
no service-policy input PM-COPP
!
interface E0/1
no service-policy input PM-COPP
!
control-plane
service-policy input PM-COPP
```

Correct Answer: *D*

HungarianDish Highly Voted 7 months, 1 week ago

Selected Answer: D

Answer "C" does not apply the policy correctly. This important part is missing from "C":
control-plane
service-policy input PM-COPP

For me, "D" is the closest solution as all other options are totally wrong. However, "D" does not limit SSH traffic to desired CIR because the drop action is missing (from the exceed or violate parameters.)

upvoted 5 times

inteldarvid Most Recent 5 months, 1 week ago

Selected Answer: D

D is correct :)

upvoted 1 times

6dd4aa0 8 months, 2 weeks ago

Selected Answer: C

SSH traffic needs to be configured such that the CIR must be policed with certain rate. In this way, SSH traffic can be controlled by the service-policy.

Hence, the answer is C

upvoted 3 times

ellen_AA 11 months ago

Selected Answer: D

D is correct!

upvoted 4 times

Question #385 Topic 1

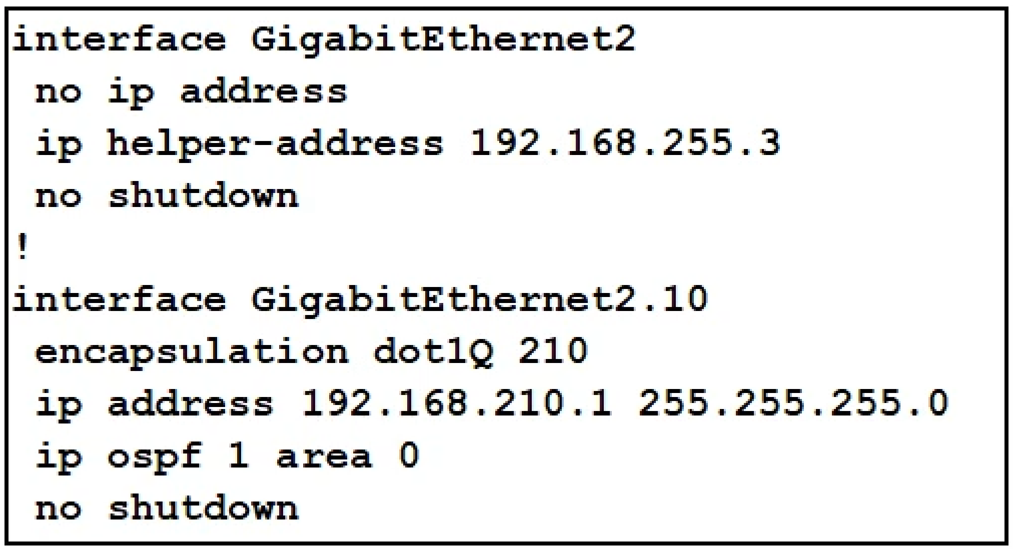

```
interface GigabitEthernet2
 no ip address
 ip helper-address 192.168.255.3
 no shutdown
!
interface GigabitEthernet2.10
 encapsulation dot1Q 210
 ip address 192.168.210.1 255.255.255.0
 ip ospf 1 area 0
 no shutdown
```

Refer to the exhibit. With the partial configuration of a router-on-a-stick, clients in VLAN 10 on Gi2 cannot obtain IP configuration from the central DHCP server. The DHCP server is reachable by a successful ping from the router. Which action resolves the issue?

A. Configure the ip helper-address 192.168.255.3 command on the Gi2.10 subinterface.

B. Configure a valid IP address on the Gi2 interface so that DHCP requests can be forwarded.

C. Configure the ip dhcp pool 1 and network 192.168.210.0 255.255.255.0 commands.

D. Configure the ip dhcp excluded-address 192.168.255.3 command on the Gi2.10 subinterface.

Correct Answer: *A*

Slinky 7 months, 2 weeks ago

Selected Answer: A

A is the correct answer, but this won't work without fixing the encapsulation to be "encapsulation dot1q 10"

upvoted 4 times

forccnp 9 months, 1 week ago

Selected Answer: A

A is the correct answer

upvoted 1 times

Question #386 *Topic 1*

The IPv6 network is under attack by an unknown source that is neither in the binding table nor learned through neighbor discovery. Which feature helps prevent the attack?

A. IPv6 Destination Guard

B. IPv6 Prefix Guard

C. IPv6 Router Advertisement Guard

D. IPv6 Snooping

Correct Answer: *B*

siyamak 4 months ago

The correct answer is D.
IPv6 Snooping learns and secures bindings for stateless autoconfiguration addresses in Layer 2 neighbor tables and analyzes ND messages in order to build a trusted binding table. IPv6 ND messages that do not have valid bindings are dropped.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: B

100 % option B :
team please look this:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/15-e/ip6f-15-e-book/ip6f-15-e-book_chapter_0110.pdf

upvoted 2 times

steve_lee 6 months, 3 weeks ago

Selected Answer: D

I would vote answer D.
To protect unknown source and ND attack.
Cisco Document (Page 2): https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-16-10/ip6f-xe-16-10-book/ip6-snooping.pdf

upvoted 2 times

HungarianDish 7 months, 1 week ago

Selected Answer: B

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-16/ip6f-xe-16-book/ip6-src-guard.pdf
IPv6 Prefix Guard prevents home-node sourcing traffic outside of the authorized and delegated traffic.
...often used when IPv6 prefixes are delegated to devices using DHCP prefix delegation.
The feature discovers ranges of addresses assigned to the link and blocks any traffic sourced with an address outside this range.

Not "A".
https://www.ciscolive.com/c/dam/r/ciscolive/global-event/docs/2022/pdf/BRKENT-3002.pdf
Destination Guard
Drops packets for destinations without a binding entry

upvoted 2 times

Malasxd 7 months, 2 weeks ago

Selected Answer: B

B is correct. It's says the SOURCE is unkown. The destination is known, so it is in binding table and the destination guard won't works to prevent it.

upvoted 1 times

sasasan12345 9 months, 1 week ago

Selected Answer: B

B is correct.The IPv6 Prefix Guard feature works within the IPv6 Source Guard feature and enables a device to reject traffic originating from addresses that are topologically incorrect.

upvoted 2 times

Lilienen 10 months ago

Selected Answer: A

IPv6 Destination Guard

upvoted 1 times

Question #387 *Topic 1*

The network administrator configured CoPP so that all routing protocol traffic toward the router CPU is limited to 1 mbps. All traffic that exceeds this limit must be dropped. The router is running BGP and OSPF. Management traffic for Telnet and SSH must be limited to 500 kbps.

```
access-list 100 permit tcp any any eq 179
access-list 100 permit tcp any any range 22 23
access-list 100 permit ospf any any
!
class-map CM-ROUTING
match access-group 100
class-map CM-MGMT
match access-group 100
!
policy-map PM-COPP
class CM-ROUTING
police 1000000 conform-action transmit
class CM-MGMT
police 500000 conform-action transmit
!
control-plane
service-policy output PM-COPP
```

No traffic is filtering through CoPP, which is resulting in high CPU utilization. Which configuration resolves the issue?

A. control-plane
```
no service-policy output PM-COPP
service-policy input PM-COPP
```

B. no access-list 100
```
access-list 100 permit tcp any any eq 179
access-list 100 permit ospf any any
access-list 101 permit tcp any any range 22 23
!
!
class-map CM-MGMT
no match access-group 100
match access-group 101
```

C. no access-list 100
```
access-list 100 permit tcp any any eq 179
access-list 100 permit ospf any any
access-list 101 permit tcp any any range 22 23
!
!
class-map CM-MGMT
no match access-group 100
match access-group 101
!
control-plane
no service-policy output PM-COPP
service-policy input PM-COPP
```

D. No access-list 100 -
```
access-list 100 permit tcp any any eq 179
access-list 100 permit tcp any any range eq 22
access-list 100 permit tcp any any range eq 23
access-list 100 permit ospf any any
```

Correct Answer: *C*

[Removed] 4 months, 3 weeks ago

Selected Answer: C

Yes, standard access list doesn't allow range keyword, policy map is in the wrong direction

upvoted 1 times

Zizu007 11 months, 3 weeks ago

Selected Answer: C

Correct!

upvoted 4 times

Question #388 *Topic 1*

Which failure detection mechanism is used for BFD?

A. consistent rate

B. Layer 2 protocol failure

C. variable rate

D. routing protocol failure

Correct Answer: *A*

Malasxd 7 months, 2 weeks ago

Selected Answer: A

Definitely "A"

upvoted 1 times

Xerath 9 months, 3 weeks ago

Selected Answer: A

BFD provides a consistent failure detection method for network administrators. Because the network administrator can use BFD to detect forwarding path failures at a uniform rate

https://www.cisco.com/en/US/docs/ios/12_4t/ip_route/configuration/guide/t_bfd.html

upvoted 2 times

Titini 10 months, 1 week ago

Selected Answer: A

BFD uses a variable rate mechanism to achieve fast and efficient detection of failures in a network path between two routers. This is a core feature of the protocol and is widely documented in technical references and standards such as RFC 5880.

upvoted 2 times

babs 10 months, 1 week ago

Selected Answer: B

should be B ?

upvoted 1 times

Malasxd 7 months, 2 weeks ago

It asks which mechanisms the BFD use to detect failures and not where the failure is detected.

upvoted 2 times

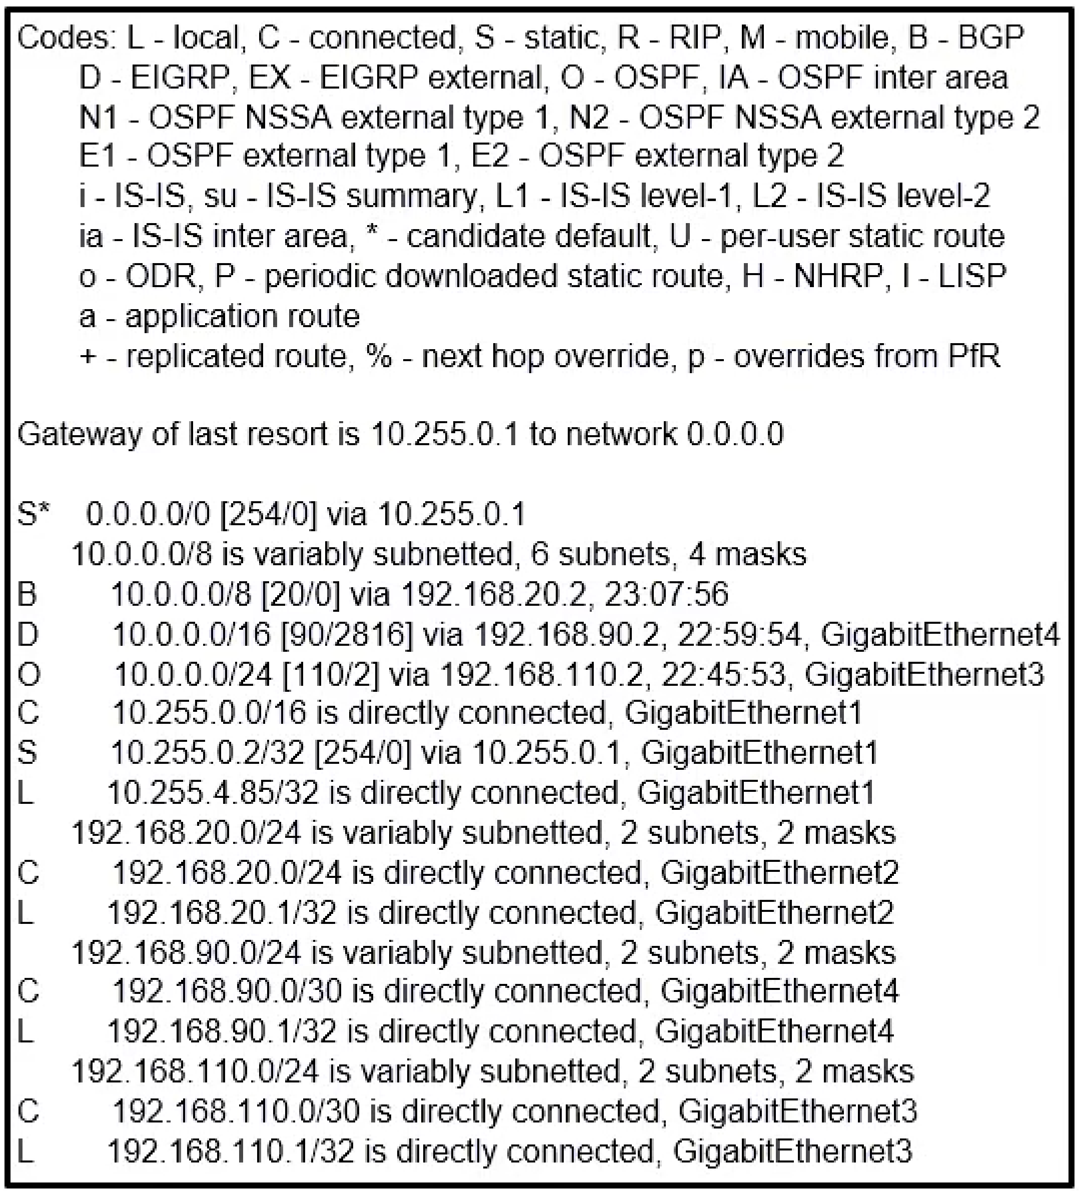

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR

Gateway of last resort is 10.255.0.1 to network 0.0.0.0

S*    0.0.0.0/0 [254/0] via 10.255.0.1
      10.0.0.0/8 is variably subnetted, 6 subnets, 4 masks
B        10.0.0.0/8 [20/0] via 192.168.20.2, 23:07:56
D        10.0.0.0/16 [90/2816] via 192.168.90.2, 22:59:54, GigabitEthernet4
O        10.0.0.0/24 [110/2] via 192.168.110.2, 22:45:53, GigabitEthernet3
C        10.255.0.0/16 is directly connected, GigabitEthernet1
S        10.255.0.2/32 [254/0] via 10.255.0.1, GigabitEthernet1
L        10.255.4.85/32 is directly connected, GigabitEthernet1
      192.168.20.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.20.0/24 is directly connected, GigabitEthernet2
L        192.168.20.1/32 is directly connected, GigabitEthernet2
      192.168.90.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.90.0/30 is directly connected, GigabitEthernet4
L        192.168.90.1/32 is directly connected, GigabitEthernet4
      192.168.110.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.110.0/30 is directly connected, GigabitEthernet3
L        192.168.110.1/32 is directly connected, GigabitEthernet3
```

Refer to the exhibit. The network administrator configured BGP as the backup route for 10.0.0.0/8 and it should work only when EIGRP 10.0.0.0/8 failed to install for site S4248T5E130F6. Which configuration resolves the issue?

A. configure terminal

!

router eigrp 1

distance eigrp 90 170

B. configure terminal

!

router eigrp 1

redistribute bgp metric 10000 1 255 1 1500

C. configure terminal

!

ip route 10.0.0.0 255.0.0.0 192.168.90.2

D. configure terminal

!

router eigrp 1

distance eigrp 10 170

Correct Answer: *D*

inteldarvid 5 months, 1 week ago

Selected Answer: D

Yes option "D"

upvoted 2 times

HungarianDish 7 months, 1 week ago

Selected Answer: D

https://www.oreilly.com/library/view/cisco-ios-in/0596008694/re197.html
distance eigrp internal-distance external-distance

AD 10 wins over AD 20 (BGP).

upvoted 2 times

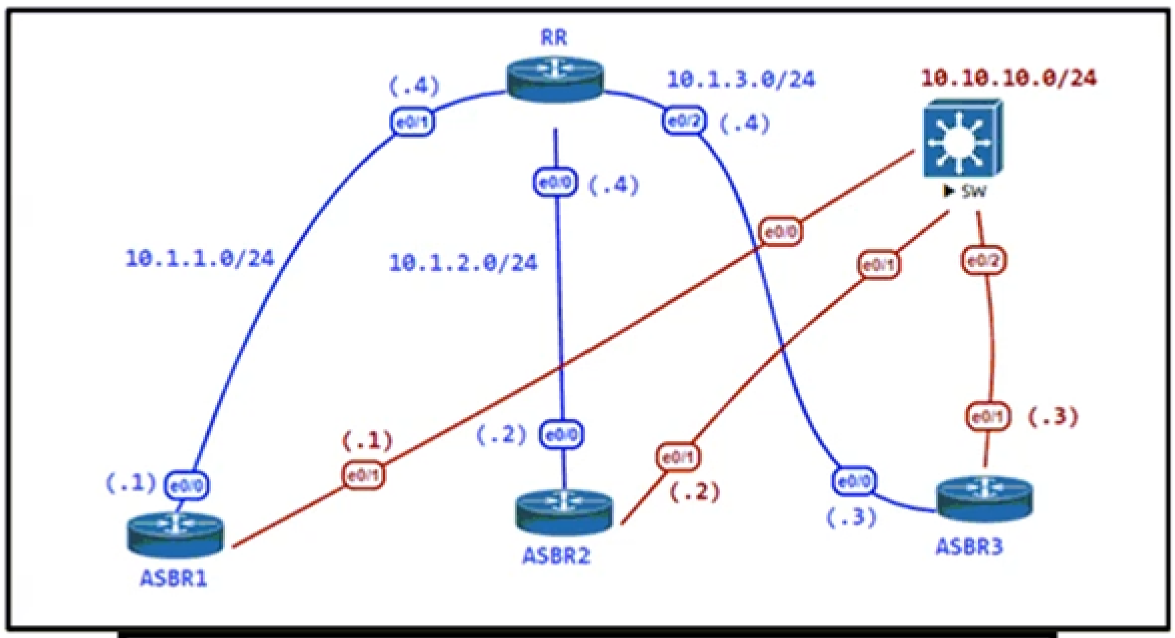

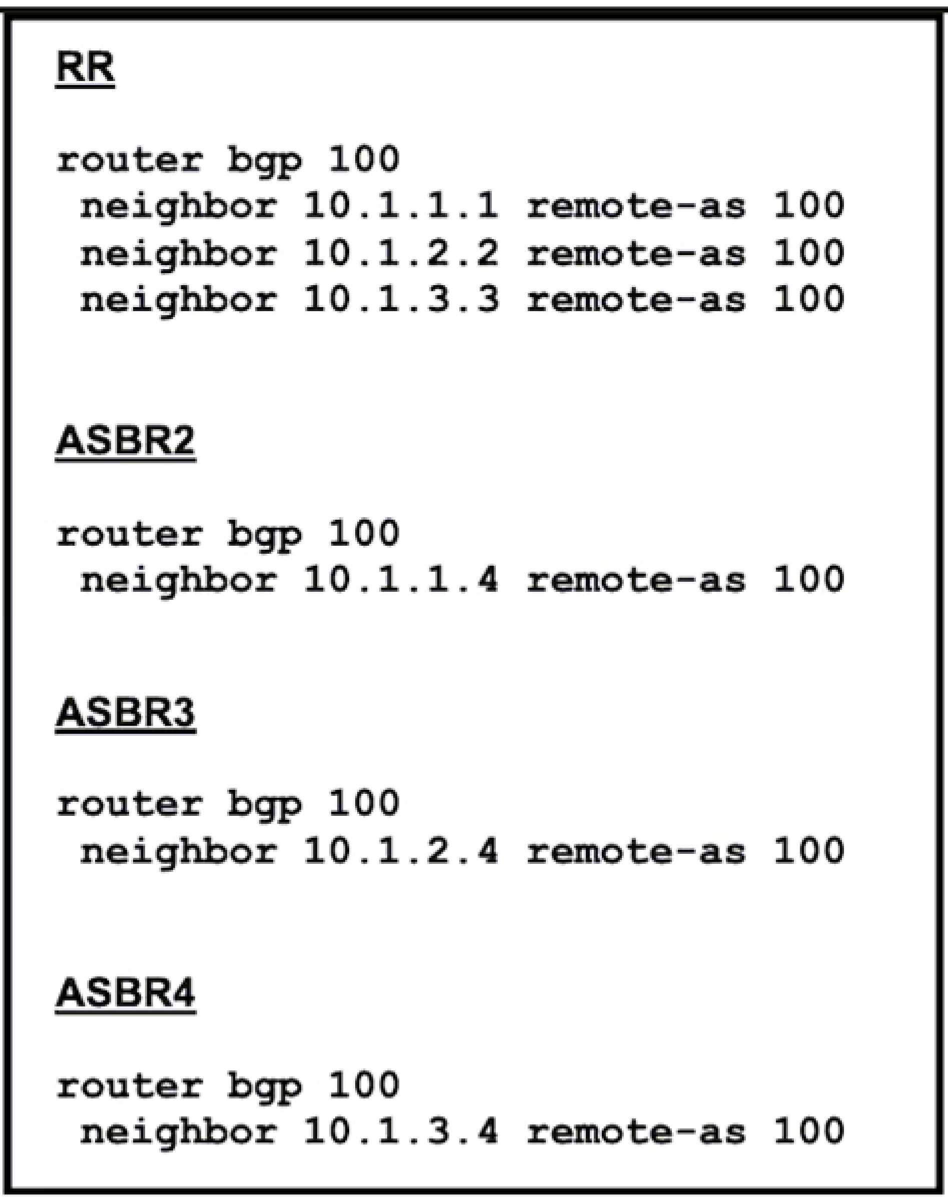

RR

```
router bgp 100
 neighbor 10.1.1.1 remote-as 100
 neighbor 10.1.2.2 remote-as 100
 neighbor 10.1.3.3 remote-as 100
```

ASBR2

```
router bgp 100
 neighbor 10.1.1.4 remote-as 100
```

ASBR3

```
router bgp 100
 neighbor 10.1.2.4 remote-as 100
```

ASBR4

```
router bgp 100
 neighbor 10.1.3.4 remote-as 100
```

Refer to the exhibit. The administrator configured the network devices for end-to-end reachability, but the ASBRs are not propagating routes to each other. Which set of configurations resolves this issue?

A. router bgp 100
neighbor 10.1.1.1 next-hop-self
neighbor 10.1.2.2 next-hop-self
neighbor 10.1.3.3 next-hop-self

B. router bgp 100
neighbor 10.1.1.1 update-source Loopback0
neighbor 10.1.2.2 update-source Loopback0
neighbor 10.1.3.3 update-source Loopback0

C. router bgp 100
neighbor 10.1.1.1 route-reflector-client
neighbor 10.1.2.2 route-reflector-client
neighbor 10.1.3.3 route-reflector-client

D. router bgp 100
neighbor 10.1.1.1 ebgp-multihop

```
neighbor 10.1.2.2 ebgp-multihop
neighbor 10.1.3.3 ebgp-muttihop
```

Correct Answer: *C*

DeWalt95 3 weeks, 4 days ago

Selected Answer: C

Its a full mesh so the best routing paths will be ensured by using Route Reflectors.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

option C correct

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: C

I think that "C" is sufficient for the solution.

upvoted 2 times

> **HungarianDish** 7 months, 1 week ago
>
> About route reflectors and next-hop-self:
> https://community.cisco.com/t5/routing/bgp-next-hop-self-not-working-while-doing-rr/td-p/2737272
>
> upvoted 1 times
>
> > **HungarianDish** 6 months, 3 weeks ago
> >
> > Confirmed solution "C" in CML lab (added and propagated a loopback on ASBR1 for testing, plus added missing network statements).
> >
> > upvoted 1 times

Slinky 7 months, 2 weeks ago

You need both rr-client and next-hop-self commands to propagate routes.

upvoted 1 times

Question #391 Topic 1

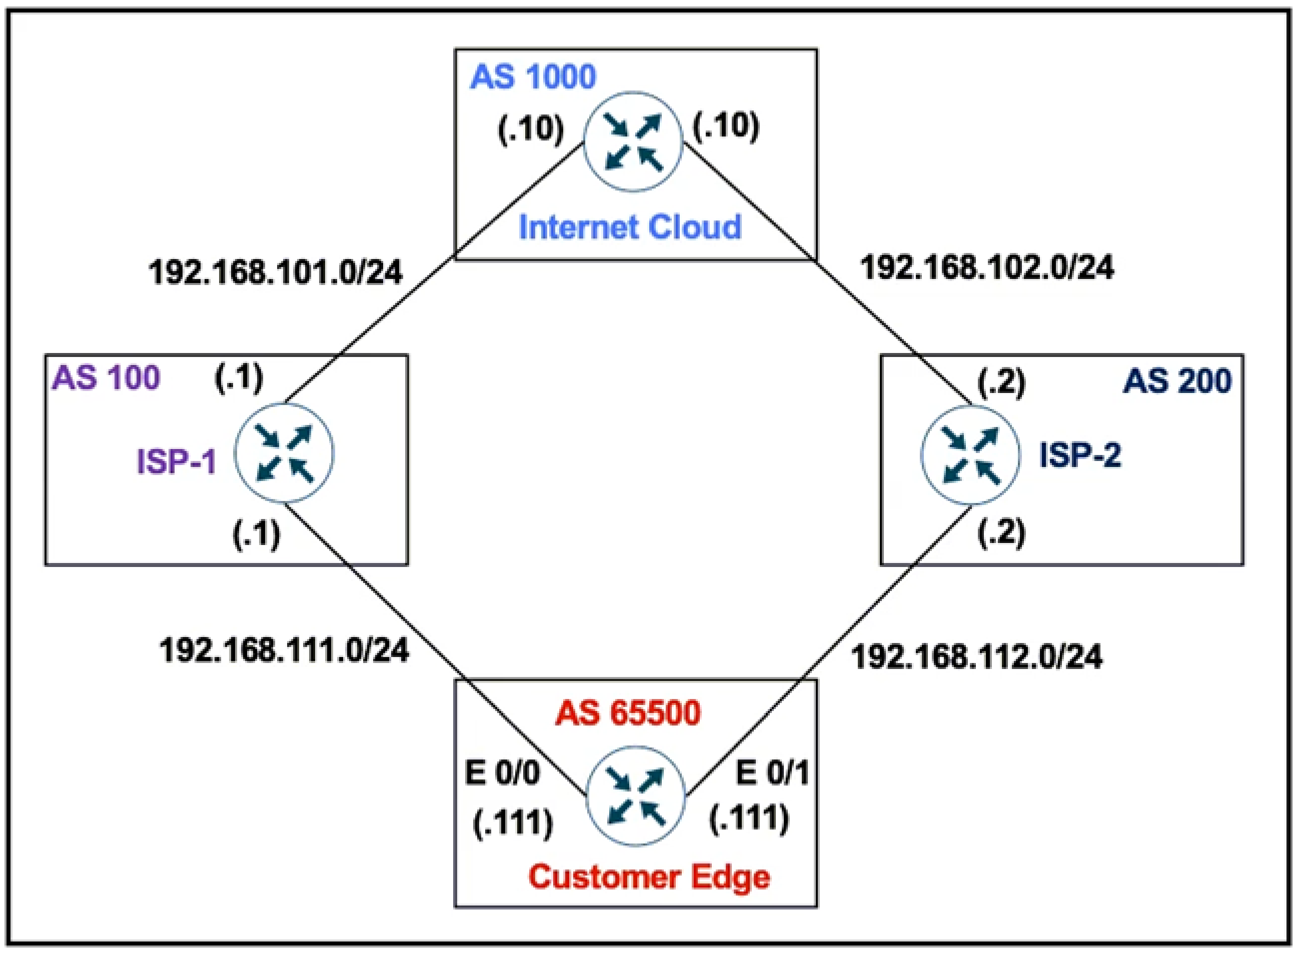

Refer to the exhibit. The Customer Edge router (AS 65500) wants to use AS 100 as the preferred ISP for all external routes.

Customer Edge -

```
route-map SETLP
set local-preference 111
!
router bgp 65500
neighbor 192.168.111.1 remote-as 100
neighbor 192.168.111.1 route-map SETLP out
neighbor 192.168.112.2 remote-as 200
```

This configuration failed to send routes to AS 100 as the preferred path. Which set of configurations resolves the issue?

A. route-map SETLP

```
set local-preference 111
!
router bgp 65500
neighbor 192.168.111.1 remote-as 100
neighbor 192.168.111.1 route-map SETLP in
```

B. route-map SETPP

```
set as-path prepend 100 100
!
router bgp 65500
neighbor 192.168.111.1 remote-as 100
neighbor 192.168.111.1 route-map SETPP in
```

C. route-map SETPP

```
set as-path prepend 111 111
!
router bgp 65500
neighbor 192.168.111.1 remote-as 100
neighbor 192.168.111.1 route-map SETPP out
```

D. route-map SETLP

```
set local-preference 111
!
```

```
router bgp 65500
neighbor 192.168.111.1 remote-as 100
neighbor 192.168.111.1 route-map SETLP out
```

Correct Answer: *A*

guy276465281819372 4 months, 3 weeks ago

Selected Answer: A

A correct

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: A

100% option A correct

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: A

We set a Local Preference of 111 to all prefixes learned from the neighbor 192.168.111.1 (AS 100), and so we make that path preferable.

https://community.cisco.com/t5/networking-knowledge-base/understanding-bgp-best-path-selection-manipulation/ta-p/3150576

upvoted 1 times

Slinky 7 months, 2 weeks ago

Selected Answer: A

You apply local pref in the inbound direction so that all externally learned prefixes are preferred from that neighbor within your AS.

upvoted 3 times

drxz 7 months, 3 weeks ago

Selected Answer: A

Answer is A.
It states that the configuration failed to send routes to AS100 as the preferred path.
You need to swap the in to out on the routemap.

with answer b you are just extending the AS path making it less desirable. which is also for incoming traffic, not outgoing.

upvoted 3 times

> **Slinky** 7 months, 2 weeks ago
>
> This is the correct answer. B is NOT it.
>
> upvoted 1 times

6dd4aa0 8 months, 2 weeks ago

A is the answer.

It cannot be B because in doing so, the AS-PATH will contain "100 100 100". This will make as though it will need additional 2 more paths to reach there!

upvoted 2 times

azzawim 8 months, 3 weeks ago

Selected Answer: B

B is the correct answer

upvoted 1 times

sasasan12345 9 months, 1 week ago

Selected Answer: A

A is correct.

upvoted 1 times

forccnp 9 months, 1 week ago

Selected Answer: B

B is the correcrt answer,
Local preference is sent to all internal BGP routers in autonomous system. Not exchanged between external BGP routers.

upvoted 1 times

mitosenoriko 11 months, 2 weeks ago

A is correct
important "in"

upvoted 2 times

Question #392 *Topic 1*

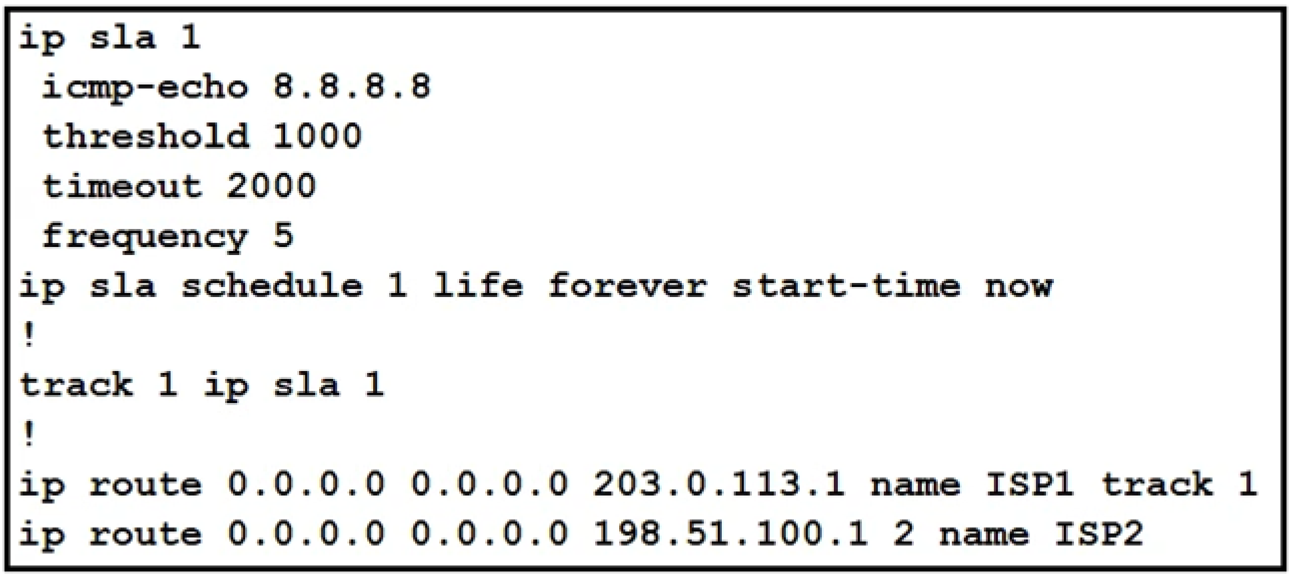

```
ip sla 1
 icmp-echo 8.8.8.8
 threshold 1000
 timeout 2000
 frequency 5
ip sla schedule 1 life forever start-time now
!
track 1 ip sla 1
!
ip route 0.0.0.0 0.0.0.0 203.0.113.1 name ISP1 track 1
ip route 0.0.0.0 0.0.0.0 198.51.100.1 2 name ISP2
```

Refer to the exhibit. The administrator noticed that the connection was flapping between the two ISPs instead of switching to ISP2 when the ISP1 failed. Which action resolves the issue?

A. Include a valid source-interface keyword in the icmp-echo statement.

B. Reference the track object 1 on the default route through ISP2 instead of ISP1.

C. Modify the static routes to refer both to the next hop and the outgoing interface.

D. Modify the threshold to match the administrative distance of the ISP2 route.

Correct Answer: *D*

DUBC89x Highly Voted 1 year ago

Selected Answer: A

If you ISP 1 fails the IP SLA will start pinging out via ISP 2. They the IP SLA will start responding again and put the static router back in for ISP 1. This will continue until ISP is back online.
IP SLA 1
icmp-echo 8.8.8.8 source-interface g1/0

upvoted 13 times

ellen_AA 11 months, 1 week ago

For more control over that, IP SLA may fall back to ISP2 in case source address can reach 8.8.8.8 by another than ISP1. You'll need to control that by adding:
ip route 8.8.8.8 255.255.255.255 g1/0
ip route 8.8.8.8 255.255.255.255 Null0 2

upvoted 2 times

HungarianDish 7 months, 1 week ago

Good point.
https://community.cisco.com/t5/routing/ip-sla-tracking-a-far-ip/td-p/1971337

upvoted 2 times

inteldarvid Most Recent 5 months, 1 week ago

Selected Answer: A

yes, option A. look this:

https://www.lead2pass.com/downloadable/download/sample/sample_id/7350/

upvoted 1 times

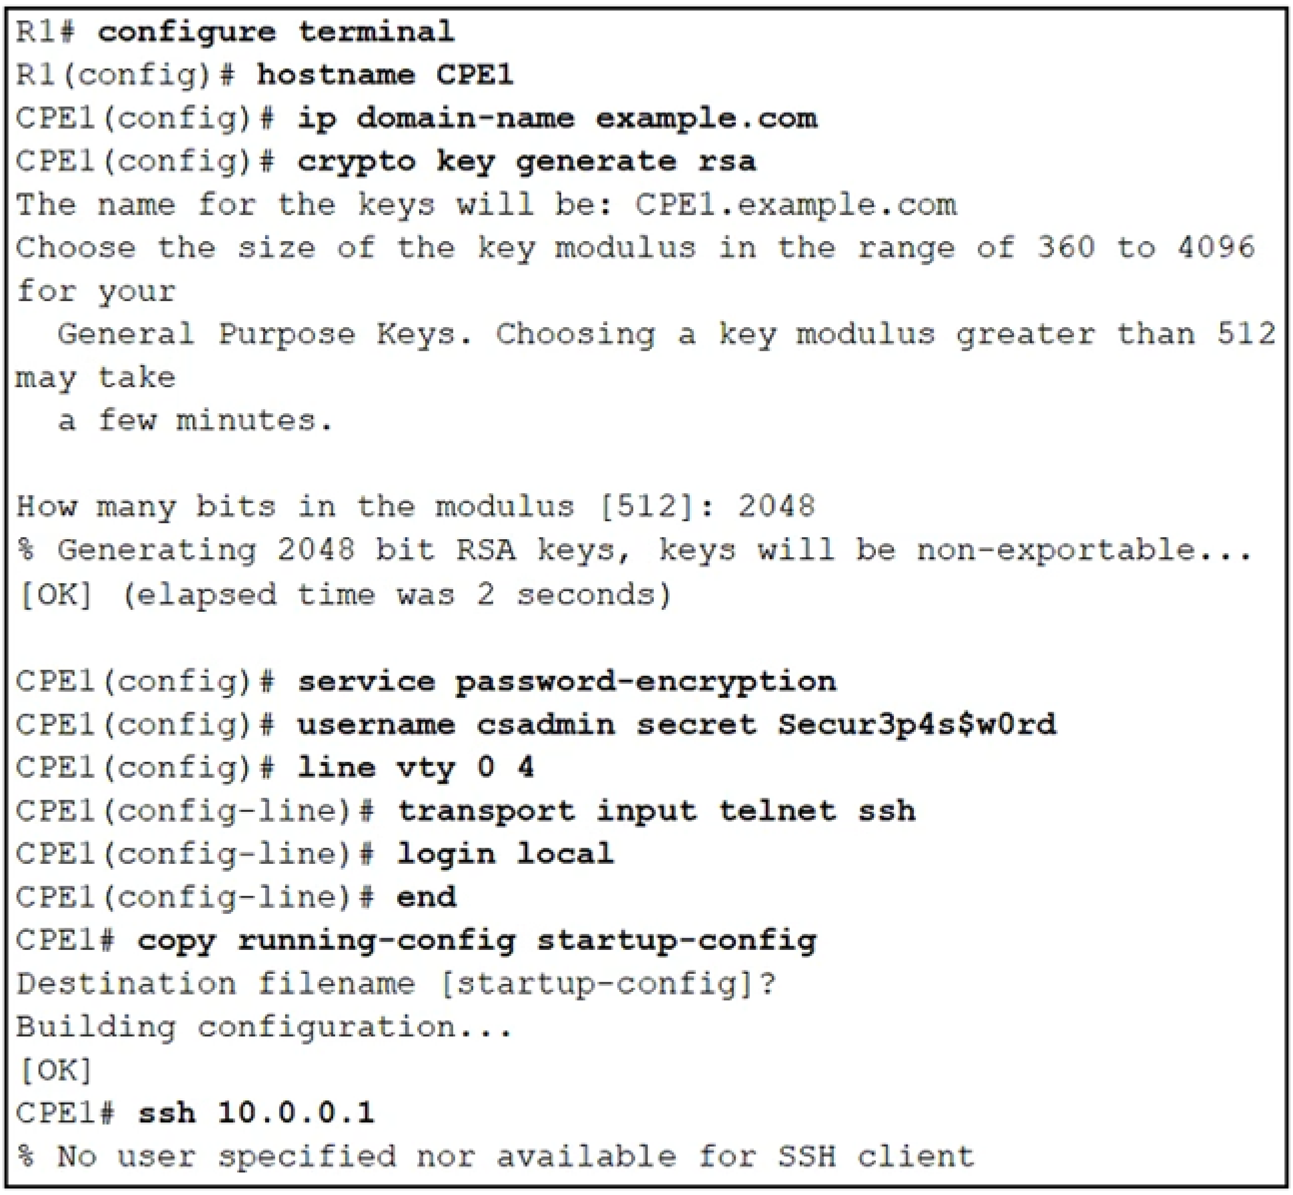

```
R1# configure terminal
R1(config)# hostname CPE1
CPE1(config)# ip domain-name example.com
CPE1(config)# crypto key generate rsa
The name for the keys will be: CPE1.example.com
Choose the size of the key modulus in the range of 360 to 4096
for your
  General Purpose Keys. Choosing a key modulus greater than 512
may take
  a few minutes.

How many bits in the modulus [512]: 2048
% Generating 2048 bit RSA keys, keys will be non-exportable...
[OK] (elapsed time was 2 seconds)

CPE1(config)# service password-encryption
CPE1(config)# username csadmin secret Secur3p4s$w0rd
CPE1(config)# line vty 0 4
CPE1(config-line)# transport input telnet ssh
CPE1(config-line)# login local
CPE1(config-line)# end
CPE1# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
[OK]
CPE1# ssh 10.0.0.1
% No user specified nor available for SSH client
```

Refer to the exhibit. An administrator must harden a router, but the administrator failed to test the SSH access successfully to the router. Which action resolves the issue?

A. SSH must be allowed with the transport output ssh command.

B. Configure enable secret to log in to the device.

C. SSH syntax must be ssh -l user ip to log in to the remote device.

D. Configure SSH on the remote device to log in using SSH.

Correct Answer: *C*

⊟ **Titini** Highly Voted 10 months, 1 week ago

Selected Answer: C

https://community.cisco.com/t5/network-management/unable-to-ssh-to-router/td-p/4047453

upvoted 5 times

⊟ **ZamanR** Most Recent 5 days, 1 hour ago

C is correct

upvoted 1 times

⊟ **fortinet1234** 2 months ago

I believe that the correct answer here is A - if you are initiating the SSH connection from the router it self you need to allow ssh out with " transport output ssh "

upvoted 1 times

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: C

correct is C

upvoted 2 times

Question #394

Topic 1

Which MPLS value is combined with the IP prefix to convert to a VPNv4 prefix?

A. 8-byte Route Distinguisher

B. 8-byte Route Target

C. 16-byte Route Target

D. 16-byte Route Distinguisher

Correct Answer: *A*

forccnp 9 months, 4 weeks ago

Selected Answer: A

Given answer is correct

upvoted 2 times

TAZZER 11 months, 1 week ago

Selected Answer: A

A is correct

upvoted 3 times

Question #395

Topic 1

What are the two reasons for RD and VPNv4 addresses in an MPLS Layer 3 VPN? (Choose two.)

A. VPN RT communities are used to identify customer unique routes.

B. When the PE redistributes customer routes into MP-BGP, they must be unique.

C. They are on a CE device to use for static configuration.

D. They are used for a BGP session with the CE device.

E. RD is prepended to each prefix to make routes unique.

Correct Answer: *BE*

inteldarvid 5 months, 1 week ago

Selected Answer: BE

Correct

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: BE

As its name implies, a route distinguisher (RD) distinguishes one set of routes (one VRF) from another. It is a unique number prepended to each route within a VRF to identify it as belonging to that particular VRF or customer. An RD is carried along with a route via MP-BGP when exchanging VPN routes with other PE routers.

https://packetlife.net/blog/2013/jun/10/route-distinguishers-and-route-targets/

upvoted 3 times

Question #396

Topic 1

An engineer configured a leak-map command to summarize EIGRP routes and advertise specifically loopback 0 with an IP of 10.1.1.1 255.255.255.252 along with the summary route. After finishing configuration, the customer complained about not receiving the summary route with the specific loopback address. Which two configurations will fix this issue? (Choose two.)

```
router eigrp 1
!
route-map Leak-Route deny 10
!
interface Serial 0/0
ip summary-address eigrp 1 10.0.0.0 255.0.0.0 leak-map Leak-Route
```

A. access-list 1 permit 10.1.1.0 0.0.0.3

B. access-list 1 permit 10.1.1.1 0.0.0.252

C. access-list 1 and match under route-map Leak-Route

D. route-map Leak-Route permit 10 and match access-list 1

E. route-map Leak-Route permit 20

Correct Answer: *AD*

inteldarvid 5 months, 1 week ago

Selected Answer: AD

A y D correct

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: AD

https://networklessons.com/cisco/ccie-routing-switching-written/eigrp-summary-leak-map

upvoted 4 times

Question #397 Topic 1

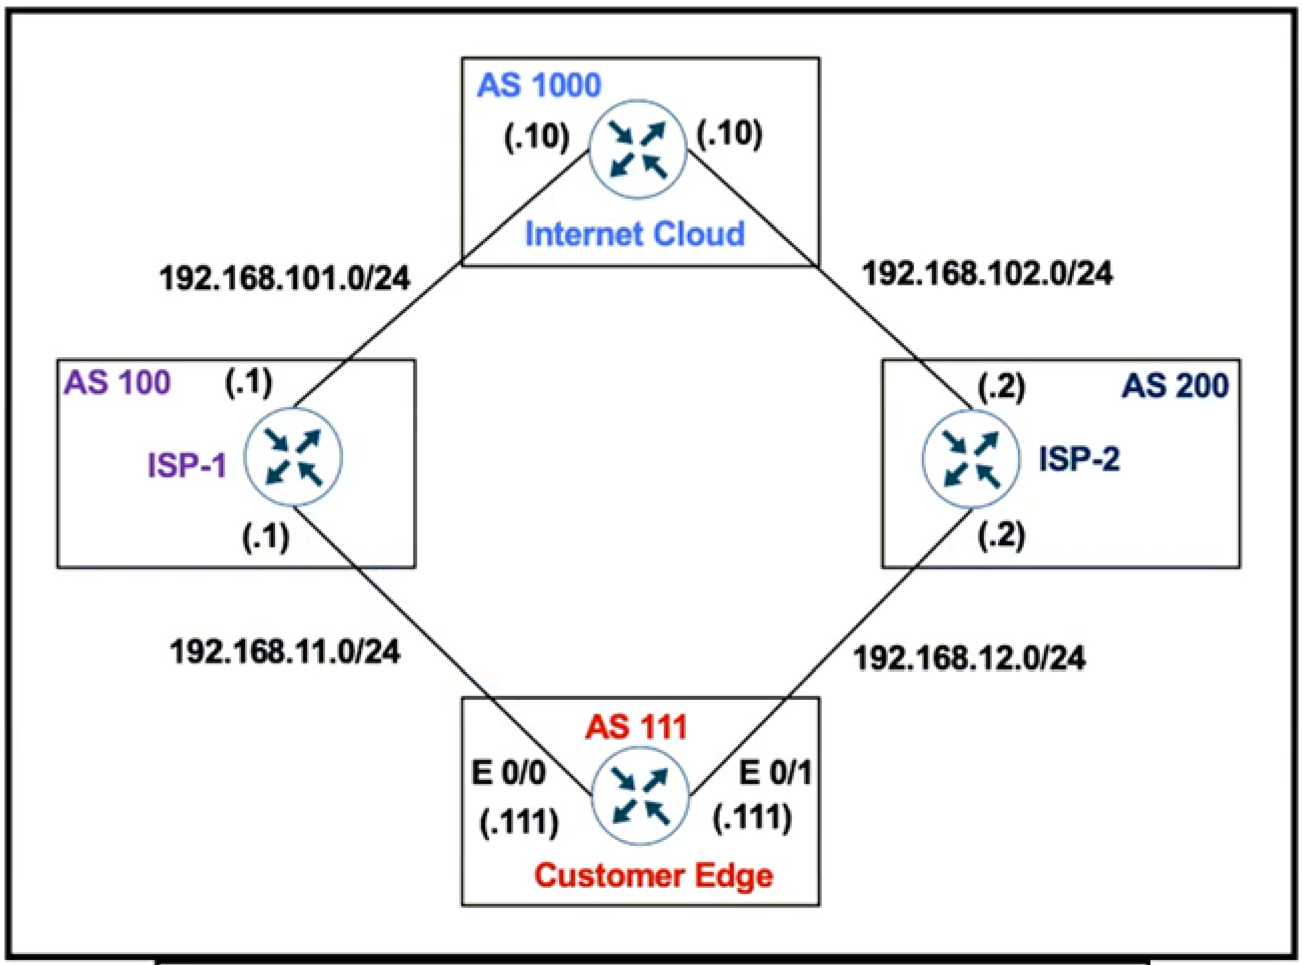

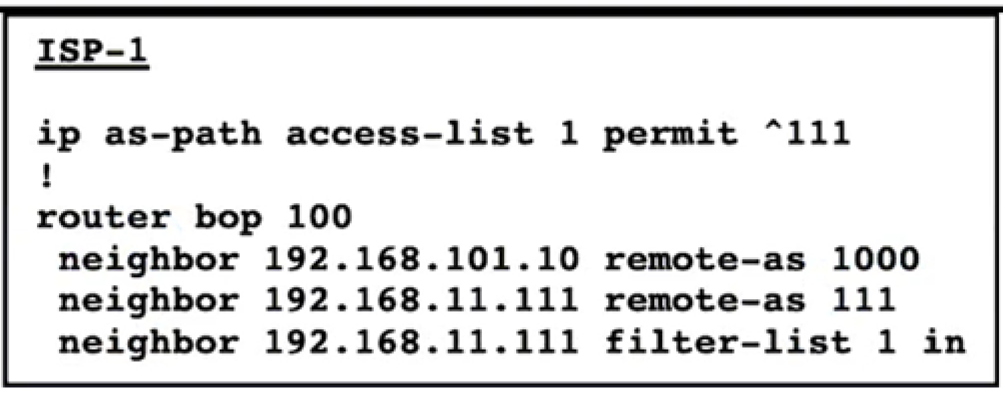

```
ISP-1

ip as-path access-list 1 permit ^111
!
router bop 100
 neighbor 192.168.101.10 remote-as 1000
 neighbor 192.168.11.111 remote-as 111
 neighbor 192.168.11.111 filter-list 1 in
```

Refer to the exhibit. AS 111 must not be used as a transit AS, but ISP-1 is getting ISP-2 routes from AS 111. Which configuration stops Customer AS from being used as a transit path on ISP-1?

A. ip as-path access-list 1 permit.*

B. ip as-path access-list 1 permit_111_

C. ip as-path access-list 1 permit ^$

D. ip as-path access-list 1 permit ^111$

Correct Answer: *C*

MJM1973 2 weeks, 6 days ago

D is the Correct Answer.
Because the question is about what can be done on the ISP-1 router so that it allows routes that originated in AS 111
ip as-path access-list 1 permit ^111$ -
^ matches beginging of the string
111 - is the string
$ matches end of the string

upvoted 1 times

aqwsdfghjklp 3 weeks ago

Why not "B"?

upvoted 1 times

Muste 4 months, 2 weeks ago

Selected Answer: D

since the router doing the configuration is the ISP the correct sintax would be *111$

upvoted 2 times

Malasxd 7 months ago

Selected Answer: C

C is right.
The local routes stills do not have it's own AS in NLRI AS-Path attribute. I am sure it's C. you can check it in BGP table, just look the local routes there and you are going to see none ASN.
upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Configuration applied on ISP-1: "ip as-path access-list 1 permit ^$" = solution "C" means receive only networks originating in the local AS (AS 100) and no Internet routes. So, no routes from AS111 are received, which is not the intended result.
upvoted 1 times

inteldarvid 5 months, 1 week ago

my friend you are worng , because the configuration is apply in ISP, not customer
upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: D

-solution C) if applied on customer edge
-solution D) if applied on ISP1 -> in this case it is

-from neighbor 192.168.11.111, receive only the routes originated from AS 111 (and no Internet routes)

At the end: clear ip bgp x.x.x.x soft in

https://community.cisco.com/t5/routing/bgp-using-as-path-filtering/td-p/1251694
https://www.ciscopress.com/articles/article.asp?p=169556
upvoted 4 times

Typovy 8 months ago

Selected Answer: C

C is correct anwer
upvoted 3 times

sasasan12345 9 months, 1 week ago

Selected Answer: D

D is correct.
upvoted 1 times

Titini 10 months, 1 week ago

Selected Answer: D

The correct configuration to prevent Customer AS 111 from being used as a transit path on ISP-1 is option D: ip as-path access-list 1 permit ^111$.

This configuration creates an access-list named "1" that permits only AS paths that consist of only AS 111. The "^" character matches the beginning of the AS path, the "$" character matches the end of the AS path, and the digits "111" match the AS number. Any other AS path, including those that pass through AS 111, will not match this access-list.
upvoted 2 times

Titini 10 months ago

Also the configuration will be applied in ISP1 not customer edge.
upvoted 1 times

ellen_AA 11 months, 1 week ago

Given answer is correct!
ISP-1 should receive from edge router (AS 111) only its locally originated route. To do that using regex, we use ^$.
upvoted 4 times

shoo83 11 months, 2 weeks ago

Answer D
Supposed to be -> ip as-path access-list 1 permit ^111$
upvoted 2 times

Hermin 10 months, 1 week ago

The question is refer to configuration stops Customer AS from being used as a transit path on "ISP-1" not on CE
upvoted 2 times

Question #398

Topic 1

What is considered the primary advantage of running BFD?

A. reduction in time needed to detect Layer 3 routing neighbor failures

B. reduction in CPU needed to detect Layer 3 routing neighbor failures

C. reduction in time needed to detect Layer 2 switched neighbor failures

D. reduction in CPU needed to detect Layer 2 switch neighbor failures

Correct Answer: *A*

HungarianDish 7 months, 1 week ago

Selected Answer: A

https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fs_bfd.html
BFD can be used at any protocol layer. However, the Cisco implementation of BFD for Cisco IOS Releases 12.2(18)SXE, 12.0(31)S, and 12.4(4)T supports only Layer 3 clients, in particular, the BGP, EIGRP, IS-IS, and OSPF routing protocols.

upvoted 3 times

Question #399 *Topic 1*

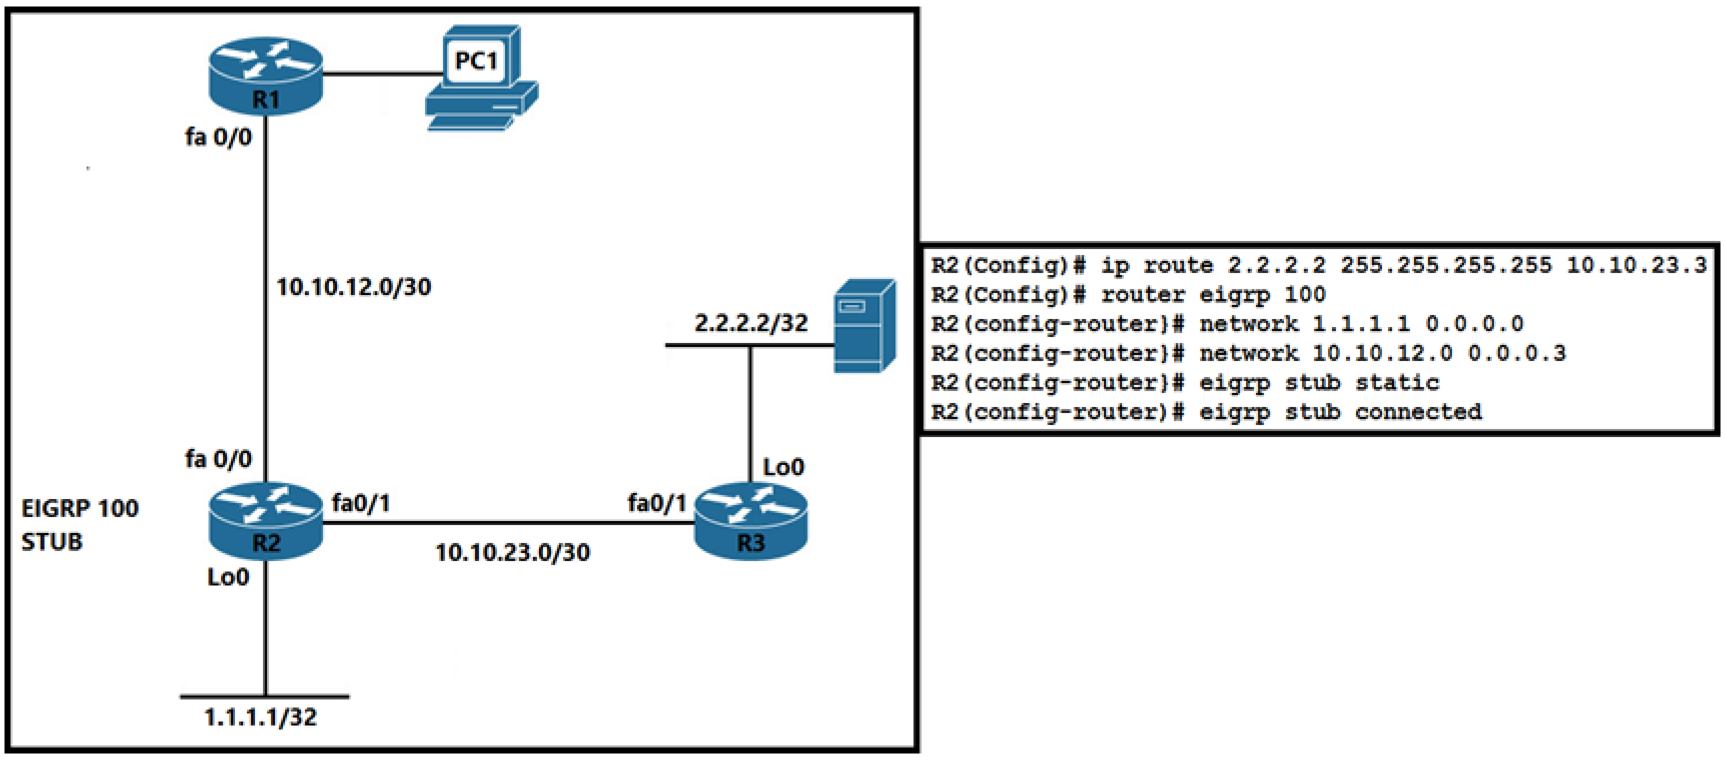

```
R2(Config)# ip route 2.2.2.2 255.255.255.255 10.10.23.3
R2(Config)# router eigrp 100
R2(config-router)# network 1.1.1.1 0.0.0.0
R2(config-router)# network 10.10.12.0 0.0.0.3
R2(config-router)# eigrp stub static
R2(config-router)# eigrp stub connected
```

Refer to the exhibit. R2 can access content on the server successfully. A network engineer finds packet drops on PC1 for traffic destined to network 2.2.2.2/32. Which action resolves the issue?

A. Redistribute the connected metric in EIGRP.

B. Add the eigrp stub connected static command.

C. Redistribute the static metric in EIGRP.

D. Remove the eigrp stub connected command.

Correct Answer: *B*

upp3r Highly Voted 8 months ago

it is true the redistribute static command is needed but after that it still won't work... as the "eigrp stub static" & "eigrp stub connected" were typed as 2 separate commands the latter overwrote the first thus still no advertising of static routes..

upvoted 11 times

HungarianDish 7 months, 1 week ago

Good point.

upvoted 1 times

MJM1973 Most Recent 2 weeks, 6 days ago

B is the correct answer from exam point of view
ommand " eigrp stub connected static summary" will advertise local connected routes, static routes and summary routes. Having said that it will advertise static routes- does not mean that it will happen automatically. You need to redistribute static routes in eigrp.

upvoted 1 times

Arsen_4 2 months, 1 week ago

I believe that C and D are needed. If you will use, as mentioned, B and C, you will resolve issue only partially. R2 will advertise server IP to R1. But what about the PC1 subnet? R3 doesn't have IP info how to reach the PC1 with EIGRP stub feature configured on R2. Unless this subnet (R1-PC1) will be configured as static on R2 it will not work. Or am I wrong?

upvoted 1 times

Arsen_4 2 months, 1 week ago

I would like to correct my answer. Only D is needed and correct answer

upvoted 2 times

fizzer 3 months, 1 week ago

Option B is right but only if the scenario I cooked up is right, otherwise question should have said select 2 answers B & C.

R2 is not advertising static route to R1 because redistribution is required
Engineer did redistribution of static and route now appears in EIGRP table but R2 is still not advertising it to R1
Realises he also needs "eigrp stub static" due to the stubbiness of R2
Added command "eigrp stub static", solved problem but introduced a new one where R2 is no longer advertising its connected Loopback to R1
Quickly added "eigrp stub connected", solved loopback advertisement issue but back to square 1

"eigrp stub connected static" is required at this point to avoid the loopback advertisement issue

upvoted 1 times

[Removed] 3 months, 2 weeks ago

Selected Answer: B

Overlapping commands would undo the redistribution of static

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

vey sorry team, my previous answer is wrong because it is not possible to add the eigrp "stub connect static command" because this applies when I have a static route with the exit interface (not with the neighbor's IP) and declare the 2.2.2.2 network on router R2 . For this exercise, the correct answer is option "C", because the stub only announces the connected networks but does not redistribute static routes. To redistribute the static route it is necessary to use the "redistribute" command. Correct option is "C". I tested it in my laboratory in GNS3

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: B

the again please veryfy, the option correct is "B". I test in my lab

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: B

yes, option B correct:

eigrp stub connected static

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

After modelling the scenario in CML: If R2 is a stub, then we need both B+C to be applied.
"ip route 2.2.2.2 255.255.255.255 10.10.23.1" is configured on R2. => Static route is installed in the RIB. Redistribution of static under eigrp process is required to get this route advertised in eigrp (to R1). = "C"
If only "eigrp stub connected" is configured, then R2 does not advertise the static route. We need "eigrp stub connected static". = "B"

upvoted 3 times

> **HungarianDish** 6 months, 3 weeks ago
>
> IP address of R3 does not fit in the ip addressing scheme on the output.
>
> upvoted 2 times
>
> > **HungarianDish** 6 months, 3 weeks ago
> >
> > 2.2.2.2/32 is received by R1 as external route if static route+redistribute static are applied. Without redistribute static (and without stub), 2.2.2.2/32 is simply advertised as internal route in the eigrp domain.
> >
> > upvoted 2 times

forccnp 9 months, 1 week ago

Selected Answer: C

Labbed it, R1 shows 2.2.2.2 even eigrp stub connected and static is configured on R2, i just redistributed static into eigrp.
C is the correct asnwer.

upvoted 2 times

shoo83 11 months, 2 weeks ago

Selected Answer: C

need to redistribute static

upvoted 4 times

> **Slinky** 7 months, 2 weeks ago
>
> That makes sense, but why do they call it "redistribute static metric?"
>
> upvoted 1 times

> **ellen_AA** 11 months ago
>
> You've got right bro: https://community.cisco.com/t5/routing/eigrp-eigrp-stub-connected-static-summary/td-p/2575321#:~:text=the%20command%20%22%20eigrp%20stub%20connected,redistribute%20static%20routes%20in%20eigrp.
>
> upvoted 2 times

heeeeyajoke 1 year ago

The stub command is will not allow R2 to advertise the nei routes to avoid being used as a transit site.

upvoted 1 times

DUBC89x 1 year ago

Selected Answer: D

Labbed it. R1 would not show an route to 2.2.2.2 until this command was removed from R2.

upvoted 4 times

Question #400

Topic 1

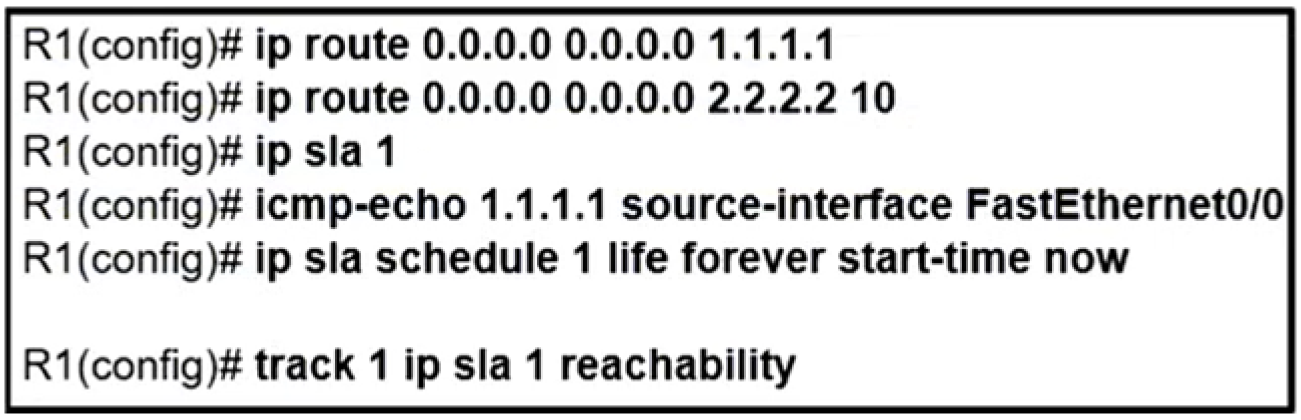

```
R1(config)# ip route 0.0.0.0 0.0.0.0 1.1.1.1
R1(config)# ip route 0.0.0.0 0.0.0.0 2.2.2.2 10
R1(config)# ip sla 1
R1(config)# icmp-echo 1.1.1.1 source-interface FastEthernet0/0
R1(config)# ip sla schedule 1 life forever start-time now

R1(config)# track 1 ip sla 1 reachability
```

Refer to the exhibit. An IP SLA is configured to use the backup default route when the primary is down, but it is not working as desired. Which command fixes the issue?

A. R1(config)# ip route 0.0.0.0 0.0.0.0 2.2.2.2 10 track 1

B. R1(config}# ip route 0.0.0.0 0.0.0.0 2.2.2.2

C. R1(config)# ip sla track 1

D. R1(config)# ip route 0.0.0.0 0.0.0.0 1.1.1.1 track 1

Correct Answer: *D*

forccnp 9 months, 1 week ago

Selected Answer: D

Given answer is correct ;)

upvoted 3 times

Question #401 Topic 1

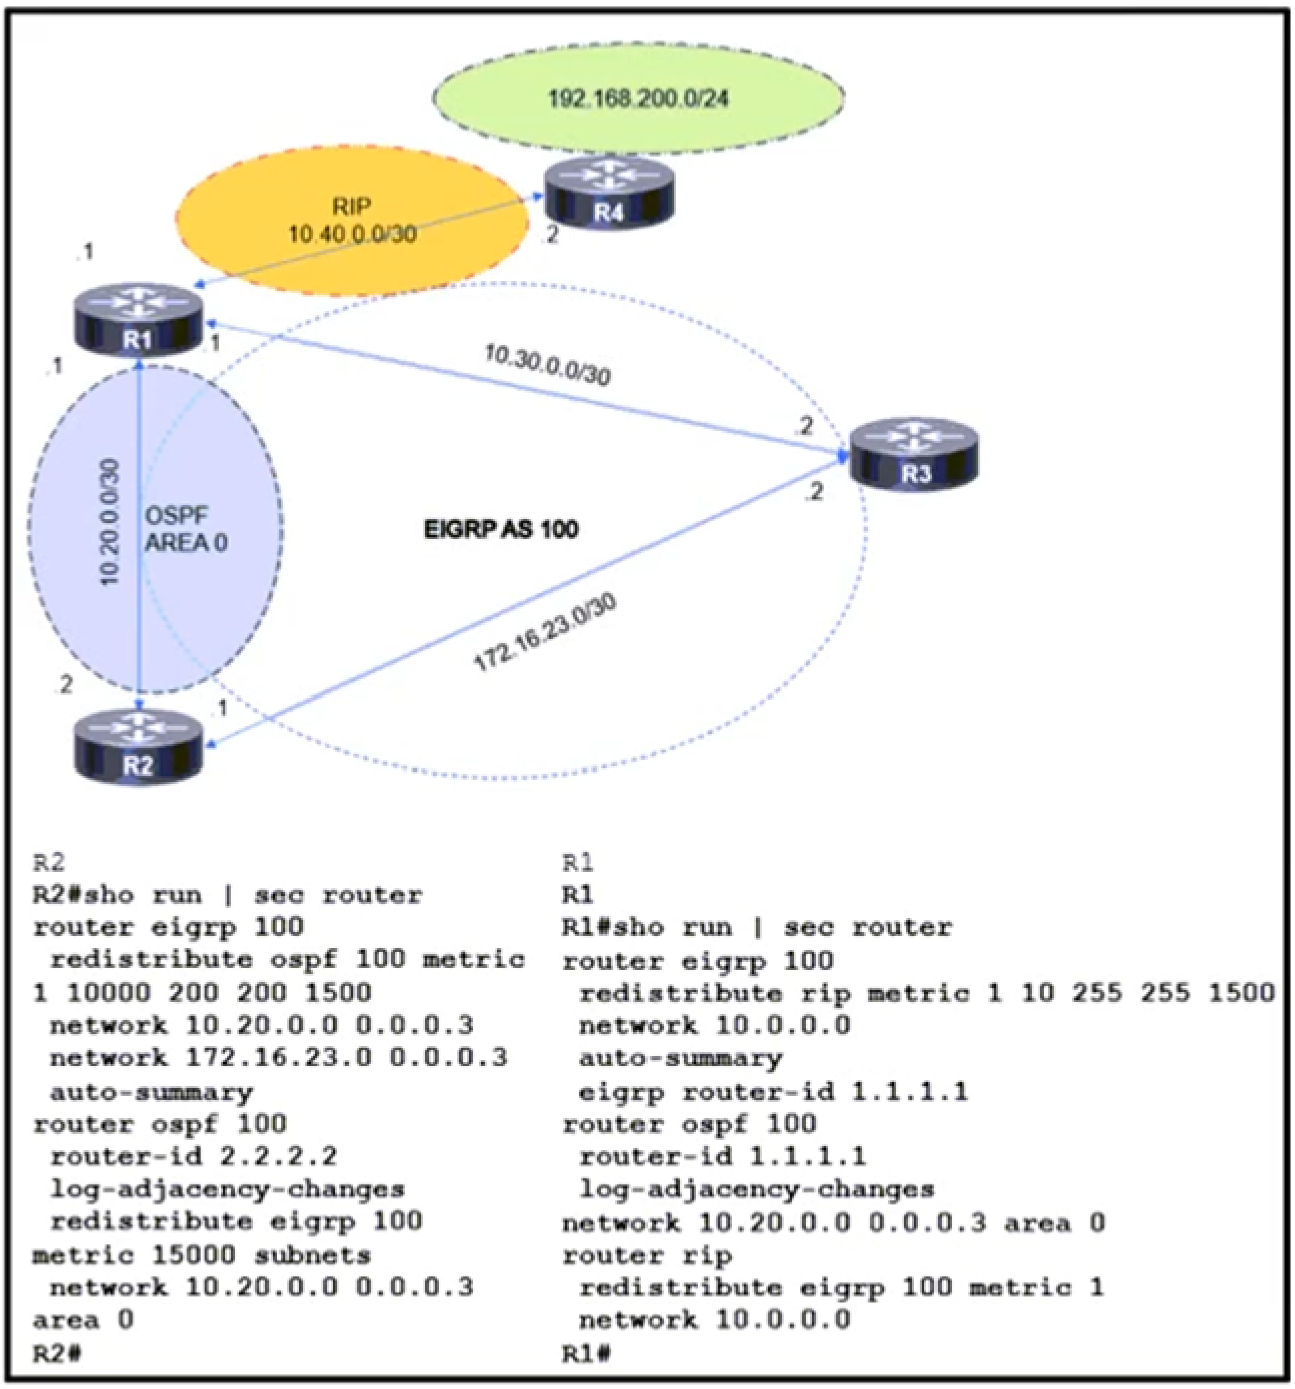

```
R2
R2#sho run | sec router
router eigrp 100
 redistribute ospf 100 metric
1 10000 200 200 1500
 network 10.20.0.0 0.0.0.3
 network 172.16.23.0 0.0.0.3
 auto-summary
router ospf 100
 router-id 2.2.2.2
 log-adjacency-changes
 redistribute eigrp 100
metric 15000 subnets
 network 10.20.0.0 0.0.0.3
area 0
R2#
```

```
R1
R1
R1#sho run | sec router
router eigrp 100
 redistribute rip metric 1 10 255 255 1500
 network 10.0.0.0
 auto-summary
 eigrp router-id 1.1.1.1
router ospf 100
 router-id 1.1.1.1
 log-adjacency-changes
network 10.20.0.0 0.0.0.3 area 0
router rip
 redistribute eigrp 100 metric 1
 network 10.0.0.0
R1#
```

Refer to the exhibit. The route to 192.168.200.0 is flapping between R1 and R2. Which set of configuration changes resolves the flapping route?

A. R2(config)#router ospf 100 -
R2(config-router)#no redistribute eigrp 100
R2(config-router)#redistribute eigrp 100 metric 1 subnets

B. R1(config)#no router rip -
R1(config)#ip route 192.168.200.0 255.255.255.0 10.40.0.2

C. R2(config)#router eigrp 100 -
R2(config-router)#no redistribute ospf 100
R2(config-router)#redistribute rip

D. R1(config)#router ospf 100 -
R1(config-router)#redistribute rip metric 1 metric-type 1 subnets

Correct Answer: *D*

inteldarvid 5 months, 1 week ago

Selected Answer: D

Option "D" correct.
the key point is router 1. because it has to redistribute rip into eigrp and rip into ospf. In this way Router 2 knows the network 192.168.200 with AD 110 and R3 with AD 170. and R1 wirh AD 120. :). Avoid fflapping :)

upvoted 2 times

Juraj22 6 months, 1 week ago

I dont understand why thet dont alsou use BGP and IS-IS in this topology......Terrible question like many others. This is not real enviroment.

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: D

Confirmed solution "D" in CML lab. See explanation from Malasxd.

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

I also labbed this scenario, and concluded that the static route of "B" only resolves the issue if we redistribute the static route on R1 (under ospf for instance). Then the prefix 192.168.200.0/24 is received by the members of the ospf and eigrp domain as an external route. Without redistributing the static route, only R1 can reach 192.168.200.0/24. No other router can.

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

With solution "D": After redistributing rip into ospf, 192.168.200.0/24 became reachable from R2, but not from R3. If R3 also needs to reach 192.168.200.0/24, then "D" is not enough.

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Sorry for this post, I tested different configurations, and made a mistake when I readded the rip configuration on R1. I spotted the error, and after applying the correct configuration "D" actually worked.

upvoted 1 times

Stylar 5 months, 1 week ago

Still flapping for me after applying D

```
R2#
O E1 192.168.200.0/24 [110/2] via 10.20.0.1, 00:01:25, GigabitEthernet2
R2#
R1#show ip route profile
IP routing table change statistics:
Frequency of changes in a 5 second sampling interval
-------------------------------------------------------------
Change/ Fwd-path Prefix Nexthop Pathcount Prefix
interval change add change change refresh
-------------------------------------------------------------
0 359 359 359 359 293
1 0 0 0 0 66
2 0 0 0 0 0

*Jul 17 18:37:47.751: RT: updating rip 192.168.200.0/24 (0x0) [local lbl/ctx:1048577/0x0] omp-tag:0 :
via 10.40.0.2 Gi1 0 0 0x0 1048578 0x100001

R 192.168.200.0/24 [120/1] via 10.40.0.2, 00:00:04, GigabitEthernet1
```

upvoted 1 times

Malasxd 7 months, 1 week ago

Selected Answer: D

D is right.

R1 just redistribute RIP in EIGRP. R2 learn 192.168.200.0 route from EIGRP and R2 redistribute EIGRP in OSPF, then R2 advertive 192.168.200.0 to R1. R1 learns 192.168.200.0 from R2 via OSPF. OSPF has AD 90 and RIP 120, so OSPF route become better than RIP route, but, for a route redistribution there is a rule that says the route must be in routing table and the source of redistribution must be the source of this route in route table, in this case, the RIP. When R1 learn this route in OSPF, the OSPF route replace the RIP route in route table, so the rule is broken and the redistribution stop working, then R1 stops redistribuing 192.168.200.0 to EIGRP and R2 stop reciving this route and R2 stops redistribute this route in OSPF so R1 won't recive this route from OSPF anymore, then OSPF route is removed from LSDB and RIB so RIP route go to the route table and the redistribution to EIGRP starts again and the problem starts over and over

If you redistribute RIP in OSPF in R1, R2 is gonna have this route as the best route from OSPF, so it does not matter if R2 learns it from EIGRP or NOT, because OSPF has AD 90 and External EIGRP 170.

upvoted 4 times

[Removed] 4 months, 3 weeks ago

Great explanation, but one correction, OSPF has AD 110, still it wins over RIP and EIGRP

upvoted 1 times

[Removed] 4 months, 3 weeks ago

External EIGRP

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Upvoted your comment. OSPF AD 110 wins over EIGRP 170.

upvoted 1 times

Typovy 9 months, 1 week ago

I labbed it and actually only B stopped route flapping but is obviously retarded solution. Can anyone share his thoughts?

upvoted 3 times

Stylar 5 months, 1 week ago

We might need to redistribute this static aswell on R1, for other routers to know about this route, otherwise they should have the default pointing to R1

upvoted 1 times

[Removed] 8 months, 1 week ago

In my understanding:
On R1 RIP is redistributed on EIGRP. Therefore it will have on EIGRP
D EX 170 192.168.200.0/24
Best route on R1 would be RIP route 120<170
On R2 EIGRP is redistributed in OSPF so R2 will have route:
O E2 110 192.168.200.0/24
Now R1 adds OSPF route on it because 110<120 and points to R2,
therefore the routing loop.

Based on alternatives B looks as the best option.

upvoted 2 times

ellen_AA 11 months, 1 week ago

This flapping would actually happen if you redistribute RIP into OSPF!!!!!!

upvoted 1 times

ellen_AA 11 months, 1 week ago

This flapping would actually happen if you redistribute RIP into EIGRP!!!!!!

upvoted 1 times

Question #402 Topic 1

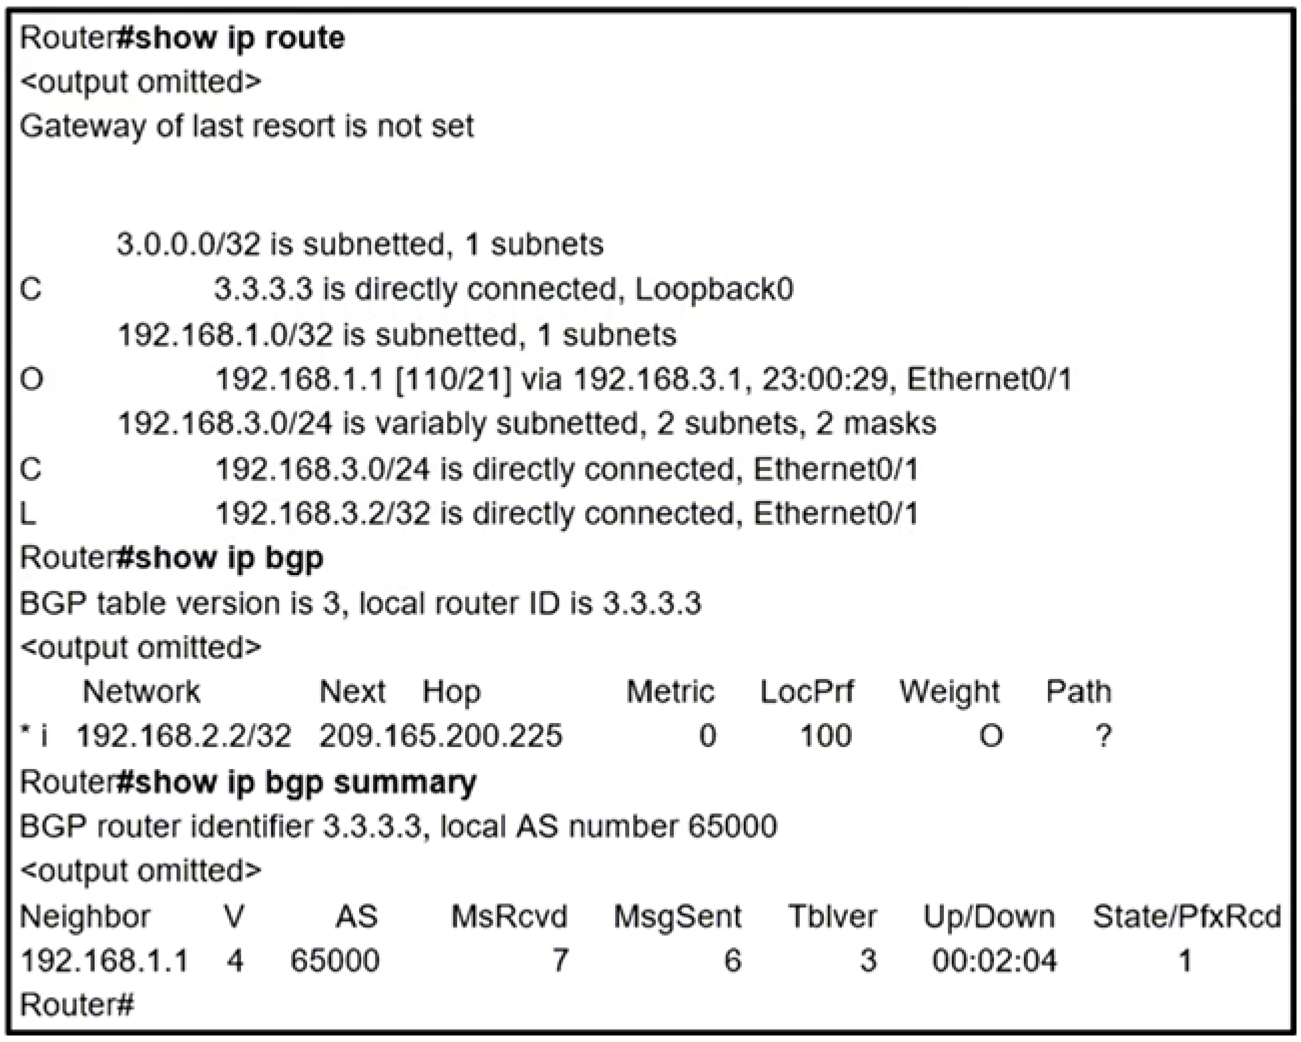

```
Router#show ip route
<output omitted>
Gateway of last resort is not set


        3.0.0.0/32 is subnetted, 1 subnets
C               3.3.3.3 is directly connected, Loopback0
        192.168.1.0/32 is subnetted, 1 subnets
O               192.168.1.1 [110/21] via 192.168.3.1, 23:00:29, Ethernet0/1
        192.168.3.0/24 is variably subnetted, 2 subnets, 2 masks
C               192.168.3.0/24 is directly connected, Ethernet0/1
L               192.168.3.2/32 is directly connected, Ethernet0/1
Router#show ip bgp
BGP table version is 3, local router ID is 3.3.3.3
<output omitted>
     Network        Next   Hop         Metric   LocPrf   Weight   Path
* i  192.168.2.2/32  209.165.200.225        0      100        O      ?
Router#show ip bgp summary
BGP router identifier 3.3.3.3, local AS number 65000
<output omitted>
Neighbor     V     AS     MsRcvd    MsgSent    Tblver    Up/Down    State/PfxRcd
192.168.1.1  4   65000         7          6         3    00:02:04          1
Router#
```

Refer to the exhibit. Which action installs route 192.168.2.2/32 in the routing table?

A. Redistribute connected networks into BGP on the local router.

B. Configure NAT on the local router to translate private IP addresses.

C. Configure the next-hop-self attribute for the peering on the local router.

D. Configure the next-hop-self attribute for the peering on the peer router.

Correct Answer: *D*

inteldarvid 5 months, 1 week ago

Selected Answer: D

yes is correct "D"

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: D

Agree on answer "D", tested the solution in CML.

upvoted 3 times

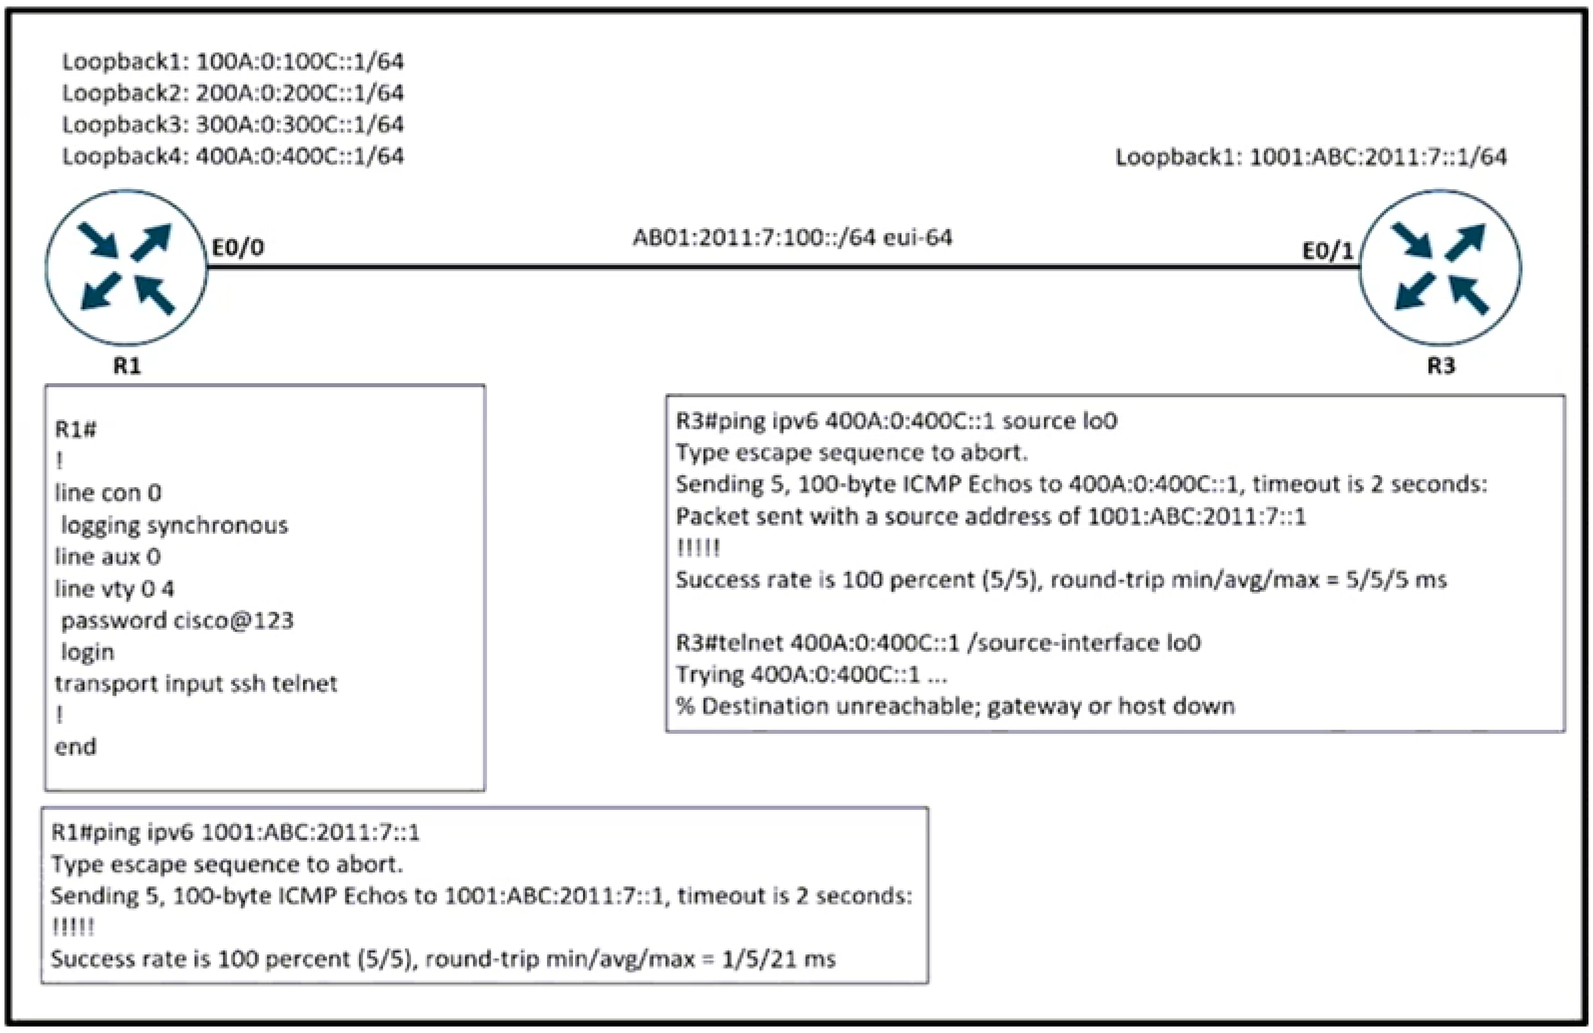
Loopback1: 100A:0:100C::1/64
Loopback2: 200A:0:200C::1/64
Loopback3: 300A:0:300C::1/64
Loopback4: 400A:0:400C::1/64
Loopback1: 1001:ABC:2011:7::1/64
E0/0
AB01:2011:7:100::/64 eui-64
E0/1
R1
R3

```
R1#
!
line con 0
 logging synchronous
line aux 0
line vty 0 4
 password cisco@123
 login
transport input ssh telnet
!
end
```

```
R3#ping ipv6 400A:0:400C::1 source lo0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 400A:0:400C::1, timeout is 2 seconds:
Packet sent with a source address of 1001:ABC:2011:7::1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/5 ms

R3#telnet 400A:0:400C::1 /source-interface lo0
Trying 400A:0:400C::1 ...
% Destination unreachable; gateway or host down
```

```
R1#ping ipv6 1001:ABC:2011:7::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1001:ABC:2011:7::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/21 ms
```

Refer to the exhibit. An engineer is trying to log in to R1 via R3 loopback address. Which action resolves the issue?

A. Add transport input SCP.

B. Remove the IPv6 traffic filter from R1, which is blocking the SSH.

C. Remove the IPv6 traffic filter from R1, which is blocking the Telnet.

D. Add transport input none.

Correct Answer: *C*

[Removed] 3 months, 2 weeks ago

Selected Answer: C

I did not see the image clearly, it has to be C but wouldn't the message say "refused by host"?

upvoted 1 times

Hummer1 3 months, 3 weeks ago

Selected Answer: C

The exhibit on the right is show ping is working and the user is using telnet to connect sourcing the loopback interface.

upvoted 1 times

[Removed] 4 months ago

Selected Answer: B

Its a fifty fifty chance between B and C. For one I would lean towards B without knowing what the hell is going, for a couple of reasons:
A trsnport input SCP is a file transfer protocol
C Telnet is least secure and hopefully is not used in production
D Transport input none would prevent remote access.

upvoted 1 times

Muste 4 months ago

you dont see the upper right image

upvoted 1 times

[Removed] 4 months ago

WHERE IS THE ACL?

upvoted 1 times

rob899 3 months, 2 weeks ago

We don't see that in the configuration. But through process of elimination, we can see that A & D are wrong because "transport input ssh telnet" already allows remote access.

So we are just left with two answers. From the image we can see the user is using "telnet" to remote access, so we must choose C. Although we cannot see an ACL, we can see all configuration looks correct and the connection is failing. So the answer is "C"

upvoted 2 times

HarwinderSekhon 4 months ago

I dont see any ACL, what am I missing?

upvoted 1 times

rob899 3 months, 2 weeks ago

We don't see that in the configuration. But through process of elimination, we can see that A & D are wrong because "transport input ssh telnet" already allows remote access.

So we are just left with two answers. From the image we can see the user is using "telnet" to remote access, so we must choose C. Although we cannot see an ACL, we can see all configuration looks correct and the connection is failing. So the answer is "C"

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

C correct

upvoted 2 times

Question #404 *Topic 1*

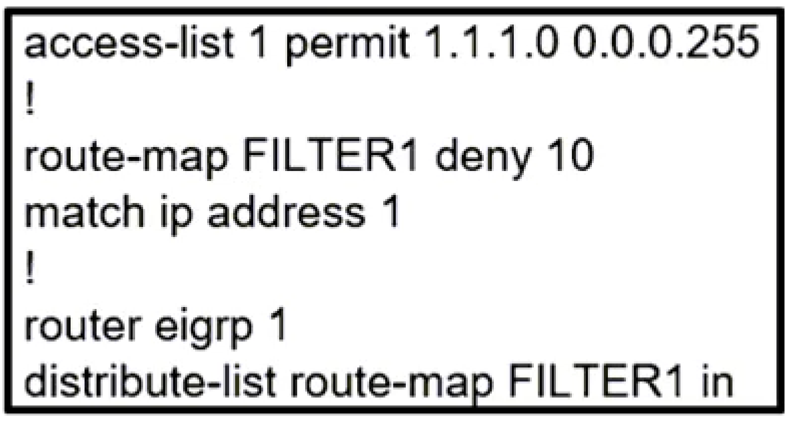

```
access-list 1 permit 1.1.1.0 0.0.0.255
!
route-map FILTER1 deny 10
match ip address 1
!
router eigrp 1
distribute-list route-map FILTER1 in
```

Refer to the exhibit. Which action restores the routes from neighbors while still filtering 1.1.1.0/24?

A. Add a second line in the access list to permit any.

B. Modify the route map to permit the access list instead of deny it.

C. Modify the access list to deny instead of permit it.

D. Add a second sequence in the route map permit 20.

Correct Answer: *D*

av3672 3 weeks, 4 days ago

Option D suggests adding a second sequence in the route map with "permit 20." This would indeed permit routes for the 1.1.1.0/24 network, but the existing deny sequence (sequence number 10) is still present. In Cisco IOS route-maps, the first match is applied. Since the deny statement comes before the permit statement in the route-map, the deny will take precedence.

Therefore, even if you add a second sequence with "permit 20," the deny sequence with "deny 10" is still in effect, and the routes for 1.1.1.0/24 would continue to be denied.

To achieve the desired result of allowing routes for 1.1.1.0/24 while still filtering other routes, you need to modify the existing deny sequence to permit the desired network. Option B correctly suggests modifying the route-map to permit the access list instead of denying it, effectively allowing routes for 1.1.1.0/24.

upvoted 1 times

av3672 3 weeks, 4 days ago

To restore the routes from neighbors while still filtering 1.1.1.0/24, you should modify the route-map to permit the routes instead of denying them. Therefore, the correct answer is:

B. Modify the route map to permit the access list instead of deny it.

upvoted 1 times

[Removed] 4 months, 3 weeks ago

Selected Answer: D

correct, route-maps have an implicit deny at end.

upvoted 1 times

Question #405

Topic 1

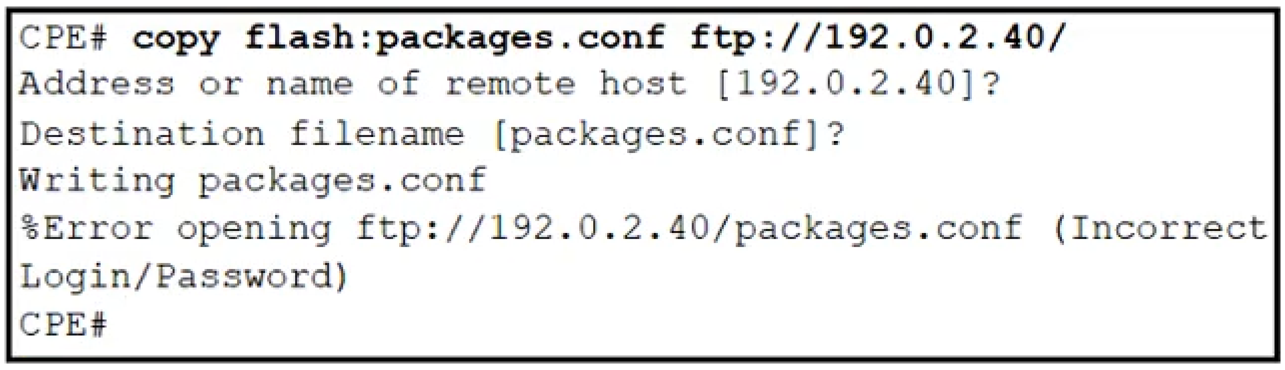

```
CPE# copy flash:packages.conf ftp://192.0.2.40/
Address or name of remote host [192.0.2.40]?
Destination filename [packages.conf]?
Writing packages.conf
%Error opening ftp://192.0.2.40/packages.conf (Incorrect
Login/Password)
CPE#
```

Refer to the exhibit. An administrator must upload the packages.conf file to an FTP server. However, the FTP server rejected anonymous service and required users to authenticate. What are the two ways to resolve the issue? (Choose two.)

A. Use the copy flash:packages.conf scp: command instead, and enter the FTP server credentials when prompted.

B. Use the copy flash:packages.conf ftp: command instead, and enter the FTP server credentials when prompted.

C. Enter the FTP server credentials directly in the FTP URL using the ftp://username:password@192.0.2.40/ syntax.

D. Create a user on the router matching the username and password on the FTP server and log in before attempting the copy.

E. Use ip ftp username and ip ftp password configuration commands to specify valid FTP server credentials.

Correct Answer: *CD*

dq28 Highly Voted 11 months, 3 weeks ago

Selected Answer: CE

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sys-image-mgmt/configuration/xe-3s/asr903/sysimgmgmt-xe-3s-asr903-book/sysimgmgmt-ftp.html

upvoted 12 times

HungarianDish Highly Voted 7 months, 1 week ago

Selected Answer: CE

https://www.oreilly.com/library/view/cisco-ios-cookbook/0596527225/ch01s15.html

upvoted 5 times

Hummer1 Most Recent 3 months, 3 weeks ago

Selected Answer: CE

Looks correct.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: CE

C and E correct

upvoted 1 times

pitcholo 10 months, 2 weeks ago

ip ftp username username Example:
Router(config)# ip ftp username user1

ip ftp password password Example:
Router(config)# ip ftp password guessme

Router# copy slot0:1:your-ios ftp://myuser:mypass@172.23.1.129/dirt/sysadmin/your-ios

upvoted 2 times

Question #406 *Topic 1*

An engineer configured VRF-Lite on a router for VRF blue and VRF red. OSPF must be enabled on each VRF to peer to a directly connected router in each VRF. Which configuration forms OSPF neighbors over the network 10.10.10.0/28 for VRF blue and 192.168.0.0/30 for VRF red?

A. router ospf 1 vrf blue
network 10.10.10.0 0.0.0.252 area 0
router ospf 2 vrf red
network 192.168.0.0 0.0.0.240 area 0

B. router ospf 1 vrf blue
network 10.10.10.0 0.0.0.15 area 0
router ospf 2 vrf red
network 192.168.0.0 0.0.0.3 area 0

C. router ospf 1 vrf blue
network 10.10.10.0 0.0.0.240 area 0
router ospf 2 vrf red
network 192.168.0.0 0.0.0.252 area 0

D. router ospf 1 vrf blue
network 10.10.10.0 0.0.0.3 area 0
router ospf 2 vrf red
network 192 168.0.0 0.0.0.15 are 0

Correct Answer: *B*

forccnp Highly Voted 9 months, 1 week ago

Selected Answer: B

B is correct answer ^_^

upvoted 5 times

ZamanR Most Recent 4 days, 23 hours ago

B is correct answer

upvoted 1 times

av3672 3 weeks, 4 days ago

Selected Answer: C

Both options cover the specified networks, but the wildcard masks in option B are incorrect. The correct wildcard masks for the specified networks are 0.0.0.240 for VRF blue and 0.0.0.252 for VRF red.

Therefore, the correct option is:

C.

upvoted 1 times

RouterToRooter 2 weeks, 6 days ago

C is wrong.
For 10.10.10.0/28
- subnet mask 255.255.255.240
- wild card 0.0.0.15
For 192.168.0.0/30
- subnet mask 255.255.255.252
- wild card 0.0.0.3

upvoted 1 times

Hummer1 5 months, 1 week ago

B is correct, D has the the wild card masks the wrong way round.

upvoted 3 times

GReddy2323 10 months, 3 weeks ago

Selected Answer: C

Wouldn't the answer be C? The network commands in B are in wildcard form.

upvoted 2 times

pitcholo 10 months, 2 weeks ago

Both B & C are correct however using wild card mask is more desirable than using the subnet mask.

Quote :

In my network implementations, I use the network statements in three different ways:

If I have to assign a specific interface into an area, I would always use network x.y.z.w 0.0.0.0 area n;
If the area address ranges are nicely assigned (which also helps immensely when you have to start summarizing), you can use a single network statement to cover the whole area. If, for example, area 3 has address range 10.1.16.0/20, use network 10.1.16.0 0.0.15.255 area 3;
If the router has all interfaces in a single area, I would almost always use network 0.0.0.0 255.255.255.255 area area-id (unless there is an extremely good reason that some interfaces should not be seen by the OSPF process).
Correct answer is B .
upvoted 2 times

Question #407 *Topic 1*

The network administrator must implement IPv6 in the network to allow only devices that not only have registered IP addresses but are also connecting from assigned locations. Which security feature must be implemented?

A. IPv6 Snooping

B. IPv6 Destination Guard

C. IPv6 Router Advertisement Guard

D. IPv6 Prefix Guard

Correct Answer: *D*

⊟ **Zizu007** 11 months, 2 weeks ago

Selected Answer: D

Correct!

https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/15S/configuration/guide/7600_15_0s_book/IPv6_Security.html#62638
upvoted 3 times

Question #408 Topic 1

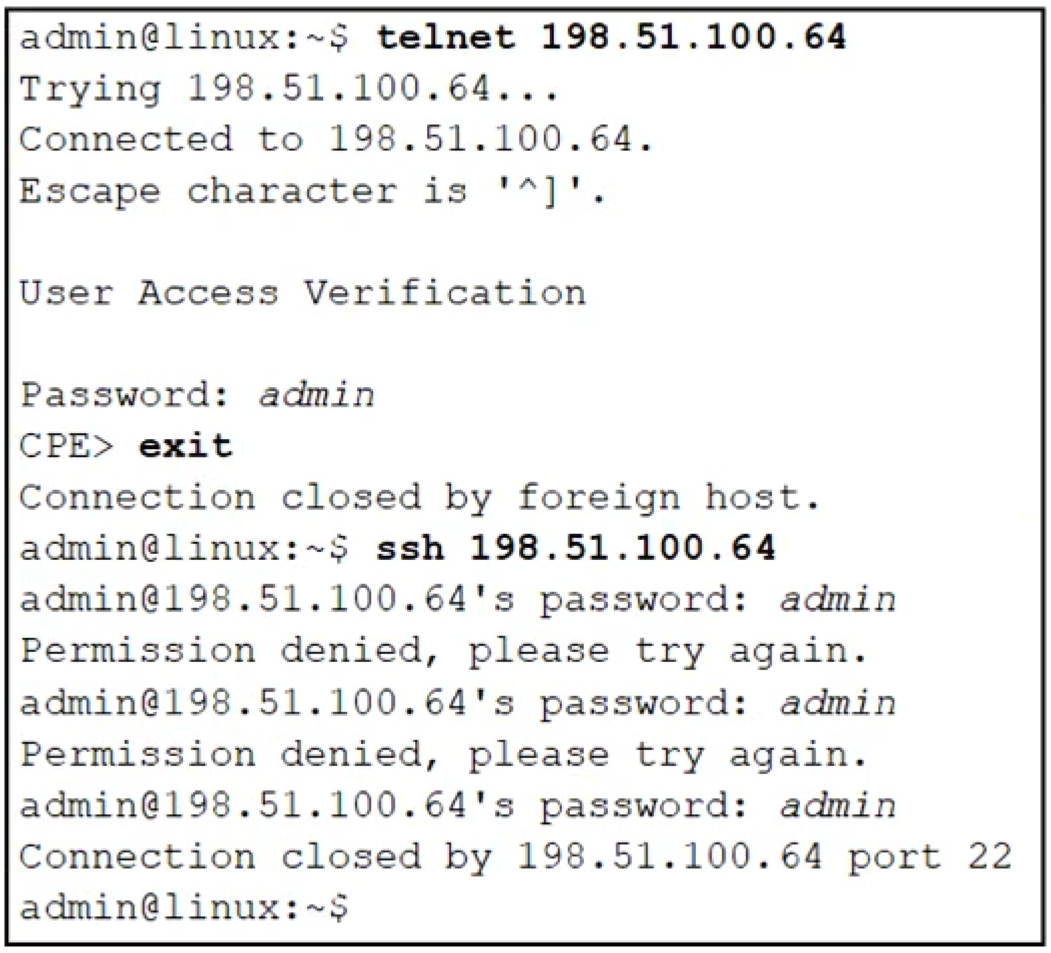

```
admin@linux:~$ telnet 198.51.100.64
Trying 198.51.100.64...
Connected to 198.51.100.64.
Escape character is '^]'.

User Access Verification

Password: admin
CPE> exit
Connection closed by foreign host.
admin@linux:~$ ssh 198.51.100.64
admin@198.51.100.64's password: admin
Permission denied, please try again.
admin@198.51.100.64's password: admin
Permission denied, please try again.
admin@198.51.100.64's password: admin
Connection closed by 198.51.100.64 port 22
admin@linux:~$
```

Refer to the exhibit. An administrator can log in to the device using Telnet, but the attempts to log in to the same device using SSH with the same credentials fail. Which action resolves this issue?

A. Configure the VTY lines with login local.

B. Configure transport input all on the VTY lines to allow SSH.

C. Configure SSH service on the router.

D. Configure to use the Telnet user database for SSH as well.

Correct Answer: *D*

inteldarvid 5 months, 1 week ago

Selected Answer: A

for mi is option "A ". I test in my lab. I need put login local in vty

upvoted 2 times

DenskyDen 6 months ago

Selected Answer: A

Labbed this. A is working.

upvoted 1 times

HungarianDish 7 months, 1 week ago

By the way, on the output we can see that the successful telnet connection did not use the local user either. No user needed to be entered for the authentication. It only used the password from the vty configuration.

upvoted 1 times

HungarianDish 7 months, 1 week ago

If we make telnet and ssh to use the local user then both username and password need to be entered for the login. For telnet, username is asked first:
cisco@PC1:~$ telnet 192.168.1.1
Connected to 192.168.1.1

Entering character mode
Escape character is '^]'.
User Access Verification

Username:

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: A

I needed to see this in my CML lab. Actually, both solutions worked fine:
A) adding "login local" to the vty configuration
and
Zizu007's solution using AAA.

The main thing was to instruct the router to use the local user database this or that way.
For this, a local user was created, e.g. username admin password cisco

As a solution with AAA is not listed, I choose answer "A".
upvoted 1 times

HungarianDish 7 months, 1 week ago

First, the same error message from the output needed to be generated to see that we catch the same issue:

MY config: transport input all + login:

```
cisco@PC1:~$ ssh -oKexAlgorithms=+diffie-hellman-group14-sha1 admin@192.168.1.1
*
admin@192.168.1.1's password:
Permission denied, please try again.
admin@192.168.1.1's password:
Permission denied, please try again.
admin@192.168.1.1's password:
Connection closed by 192.168.1.1 port 22
cisco@PC1:~$
```

(both passwords under vty config and from local user config produced this same error message)
upvoted 2 times

HungarianDish 7 months, 1 week ago

My config:

```
R1(config-line)#do sh run | sec line vty
line vty 0 4
password telnet
login
transport input all
R1(config-line)#

R1#sh run | i username
username admin password 0 cisco
R1#
```

upvoted 1 times

HungarianDish 7 months, 1 week ago

Then my config: transport input telnet + login local => different error:

```
cisco@PC1:~$ ssh -oKexAlgorithms=+diffie-hellman-group14-sha1 admin@192.168.1.1
kex_exchange_identification: Connection closed by remote host
Connection closed by 192.168.1.1 port 22
cisco@PC1:~$
```

upvoted 1 times

HungarianDish 7 months, 1 week ago

Working solutions:
- Authentication against local user database:

```
username admin password cisco
line vty 0 4
transport input telnet ssh
login local
```

or
-AAA:

```
username admin password cisco
aaa new-model
aaa authentication login default local
aaa authorization exec default local
```

Result:

```
cisco@PC1:~$ ssh -oKexAlgorithms=+diffie-hellman-group14-sha1 admin@192.168.1.1

R1>
```

upvoted 1 times

Zizu007 11 months, 2 weeks ago

Selected Answer: D

D is the best among the answers.

A - Wrong, there is mention that user is configured locally.
B - Wrong, ssh is already allowed, password prompt is presented.
C - Wrong, is already active.
D - Correct, with aaa authentication login (custom-method name) and separate VTY lines
EXAMPLE:

```
aaa new-model
aaa authentication login SSH enable
aaa session-id common
!
line vty 0 4
privilege level 15
transport input telnet
line vty 5 15
privilege level 15
login authentication SSH
transport input ssh
!
```

solution:

```
aaa authentication login SSH local
```

upvoted 3 times

Question #409 *Topic 1*

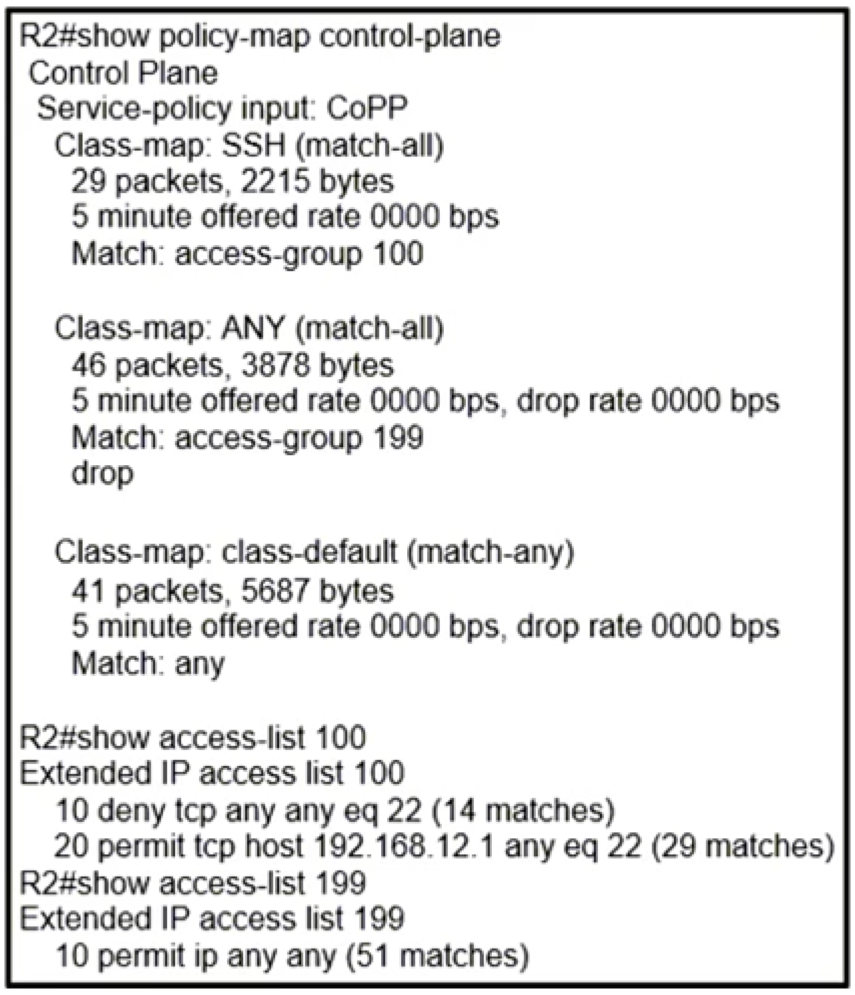

```
R2#show policy-map control-plane
 Control Plane
  Service-policy input: CoPP
    Class-map: SSH (match-all)
      29 packets, 2215 bytes
      5 minute offered rate 0000 bps
      Match: access-group 100

    Class-map: ANY (match-all)
      46 packets, 3878 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: access-group 199
      drop

    Class-map: class-default (match-any)
      41 packets, 5687 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: any

R2#show access-list 100
Extended IP access list 100
    10 deny tcp any any eq 22 (14 matches)
    20 permit tcp host 192.168.12.1 any eq 22 (29 matches)
R2#show access-list 199
Extended IP access list 199
    10 permit ip any any (51 matches)
```

Refer to the exhibit. Which action limits access to R2 from 192.168.12.1?

A. Modify sequence 20 to permit tcp host 192.168.12.1 eq 22 any to access-list 100.

B. Swap sequence 10 with sequence 20 in access-list 100.

C. Swap sequence 20 with sequence 10 in access-list 100.

D. Modify sequence 10 to deny tcp any eq 22 any to access-list 100.

Correct Answer: *B*

xzckk Highly Voted 1 year ago

What is the difference between B and C??

upvoted 16 times

Mad_Scorpion 10 months, 3 weeks ago

I guess there maybe a typo in the original question. Option C should be "swap seq 20 with seq 10 in access-list 199".

upvoted 3 times

ZamanR Most Recent 6 days, 7 hours ago

I think C is the answer

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: B

IMHO, B or C seem to be both OK. They would like to police SSH traffic only from source 192.168.12.1.
So, they need to match the traffic with an access-list:
10 permit tcp host 192.168.12.1 any eq 22 -> Police (rate-limit) this traffic
20 deny tcp any any eq 22 -> Allow this traffic unconstrained

Then they use it in the class-map SSH:
class-map SSH
match access-group 100

policy-map CoPP
class SSH

This CoPP constrains traffic from 192.168.12.1 (matched by "permit"), but allows hosts with any other source address without constraint (excluded by "deny").

They have to "deny" SSH traffic from any other source addresses in the ACL so that they are excluded from "class-map SSH".
They will be matched and allowed unconstrained by the "class-default" which is implemented implicitly at the end of the policy-map.

upvoted 2 times

HungarianDish 7 months, 1 week ago

Good example + explanation:
https://www.networktut.com/control-plane-policing-copp-tutorial

upvoted 1 times

Question #410 *Topic 1*

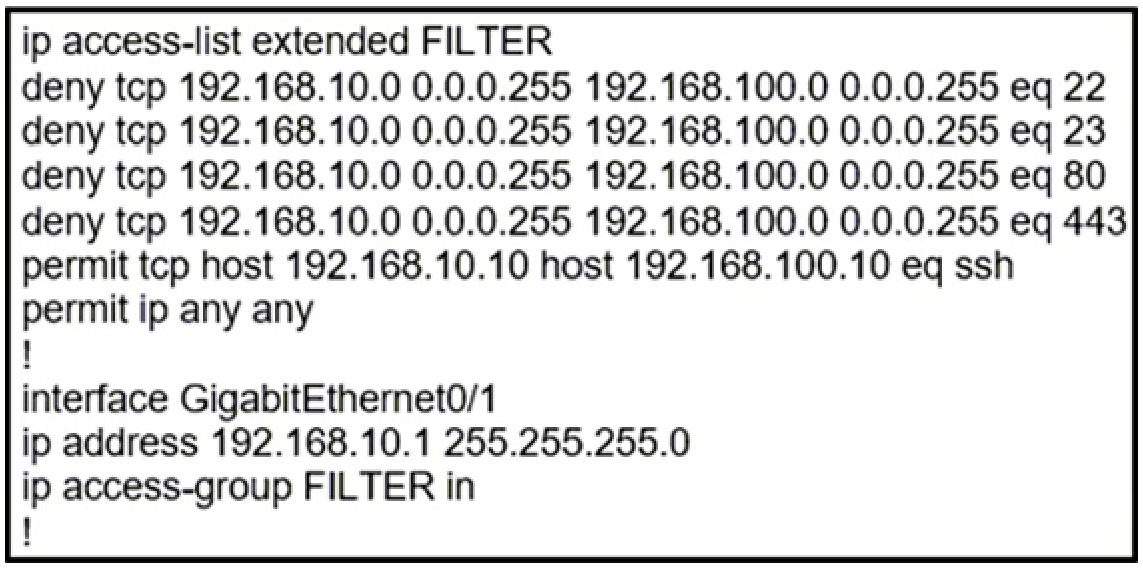

```
ip access-list extended FILTER
deny tcp 192.168.10.0 0.0.0.255 192.168.100.0 0.0.0.255 eq 22
deny tcp 192.168.10.0 0.0.0.255 192.168.100.0 0.0.0.255 eq 23
deny tcp 192.168.10.0 0.0.0.255 192.168.100.0 0.0.0.255 eq 80
deny tcp 192.168.10.0 0.0.0.255 192.168.100.0 0.0.0.255 eq 443
permit tcp host 192.168.10.10 host 192.168.100.10 eq ssh
permit ip any any
!
interface GigabitEthernet0/1
ip address 192.168.10.1 255.255.255.0
ip access-group FILTER in
!
```

Refer to the exhibit. The ACL is placed on the inbound Gigabit 0/1 interface of the router. Host 192.168.10.10 cannot SSH to host 192.168.100.10 even though the flow is permitted. Which action resolves the issue without opening full access to this router?

A. Temporarily move the permit ip any any line to the beginning of the ACL to see if the flow works

B. Temporarily remove the ACL from the interface to see if the flow works.

C. Move the SSH entry to the beginning of the ACL.

D. Run the show access-list FILTER command to view if the SSH entry has any hit statistics associated with it.

Correct Answer: *C*

inteldarvid 4 months, 3 weeks ago

Selected Answer: C

correct

upvoted 2 times

Question #411 *Topic 1*

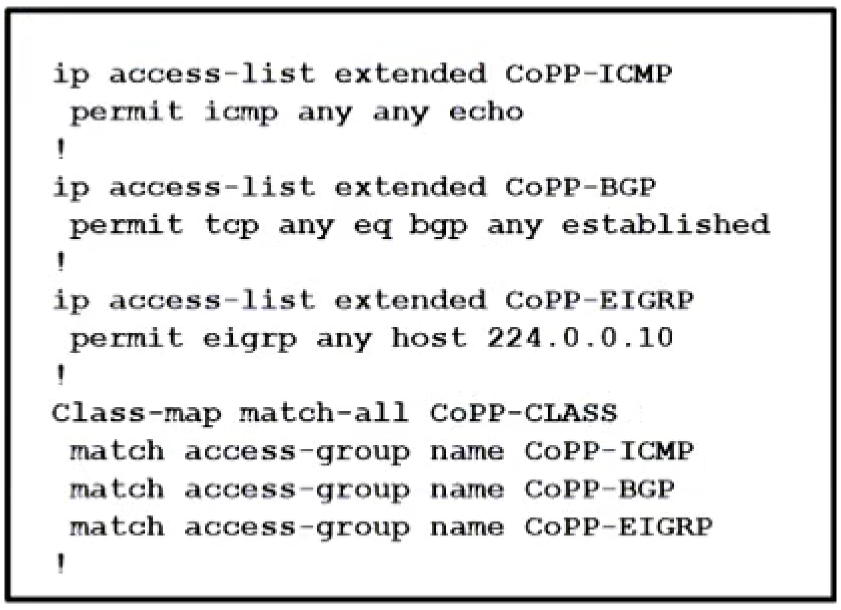

```
ip access-list extended CoPP-ICMP
 permit icmp any any echo
!
ip access-list extended CoPP-BGP
 permit tcp any eq bgp any established
!
ip access-list extended CoPP-EIGRP
 permit eigrp any host 224.0.0.10
!
Class-map match-all CoPP-CLASS
 match access-group name CoPP-ICMP
 match access-group name CoPP-BGP
 match access-group name CoPP-EIGRP
!
```

Refer to the exhibit. A CoPP policy is implemented to allow specific control traffic, but the traffic is not matching as expected and is getting unexpected behavior of control traffic. Which action resolves the issue?

A. Use match-any instruction in class-map.

B. Create a separate class map against each ACL.

C. Create a separate class map for ICMP traffic.

D. Use default-class to match ICMP traffic.

Correct Answer: *A*

ZamanR 5 days, 21 hours ago

A is correct answer

upvoted 1 times

aqwsedrfghjk 1 month ago

I don't know why B is wrong.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: A

100 % Option A

upvoted 2 times

ellen_AA 11 months, 1 week ago

Given answer is correct!
-match-all in class-map statement means that all ACLs have to match (which is very unlikely).
-match-any in class-map statement means that at least one ACL has to match. For this reason why no traffic was matching.

upvoted 4 times

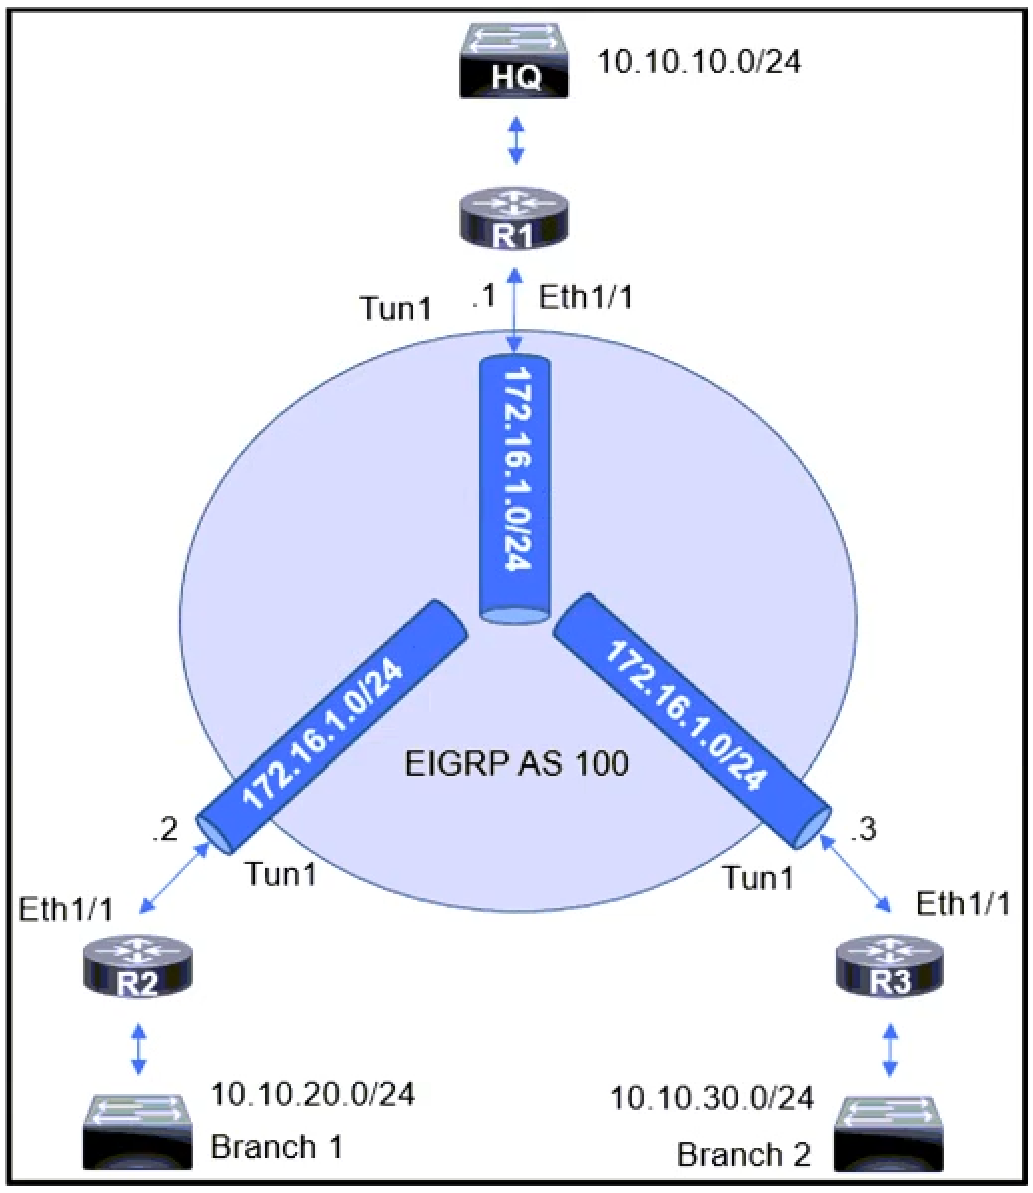

Refer to the exhibit. An engineer sets up a DMVPN connection to connect branch 1 and branch 2 to HQ. Branch 1 and branch 2 cannot communicate with each other. Which change must be made to resolve this issue?

A. R1(config)#int eth1/1 -
R1(config-if)#no ip split-horizon eigrp 100

B. R2(config)#router eigrp 100 -
R2(config-router)#tneighbor 172.16.1.3

C. R3(config)#router eigrp 100 -
R3(config-router)#neighbor 172.16.1.2

D. R1(config)#int tunnel 1 -
R1(config-if)#no ip split-horizon eigrp 100

Correct Answer: *D*

JieW 3 months, 4 weeks ago

Selected Answer: D

answer correct

upvoted 2 times

Question #413 *Topic 1*

The network administrator is tasked to configure R1 to authenticate telnet connections based on Cisco ISE using RADIUS. ISE has been configured with an IP address of 192.168.1.5 and with a network device pointing towards R1 (192.168.1.1) with a shared secret password of Cisco123. If ISE is down, the administrator should be able to connect using the local database with a username and password combination of admin/cisco123.

The administrator has configured the following on R1:

```
aaa new-model
!
username admin password cisco123
!
radius server ISE1
address ipv4 192.168.1.5
key Cisco123
!
aaa group server tacacs+ RAD-SERV
server name ISE1
!
aaa authentication login RAD-LOCAL group RAD-SERV
```

ISE has gone down. The Network Administrator is not able to Telnet to R1 when ISE went down. Which two configuration changes will fix the issue? (Choose two.)

A. aaa authentication login RAD-SERV group RAD-LOCAL local

B. aaa authentication login RAD-LOCAL group RAD-SERV local

C. line vty 0 4
login authentication RAD-LOCAL

D. line vty 0 4
login authentication default

E. line vty 0 4
login authentication RAD-SERV

Correct Answer: *BC*

⊟ **ellen_AA** Highly Voted 11 months, 1 week ago

Given answers are correct;
One more detail, we are asked to configure radius not tacas+.
So, configure:
- aaa group server radius RAD-SERV
instead of:
- aaa group server tacas+ RAD-SERV

upvoted 6 times

⊟ **HarwinderSekhon** Most Recent 4 months ago

just know that "RAD-LOCAL" is name of login method. like aaa authentication "default".

upvoted 2 times

⊟ **inteldarvid** 4 months, 3 weeks ago

Selected Answer: BC

B and C correct

upvoted 1 times

Question #414 *Topic 1*

The network administrator deployed the Binding Table Recovery feature. Which two devices recover the missing binding table entries? (Choose two.)

A. DHCP client

B. destination host

C. DHCP relay agent

D. source host

E. DHCP server

Correct Answer: *BE*

inteldarvid 5 months, 1 week ago

Selected Answer: BE

yes, B and E:
This question is very crazy. WTF cisco. We have memorize all book

This mechanism enables the binding table to recover in the event of a device reboot. The recovery mechanism will block any data traffic sourced from an unknown source; that is, a source not already specified in the binding table and previously learned through ND or DHCPgleaning. Thisfeature recoversthe missing binding table entries when the resolution for a destination address fails in the destination guard. When a failure occurs, a binding table entry is recovered by querying the DHCP server or the destination host, depending on the configuration

upvoted 2 times

DUBC89x 1 year ago

Selected Answer: BE

This mechanism enables the binding table to recover in the event of a device reboot. The recovery mechanism will block any data traffic sourced from an unknown source; that is, a source not already specified in the binding table and previously learned through ND or DHCPgleaning. Thisfeature recoversthe missing binding table entries when the resolution for a destination address fails in the destination guard. When a failure occurs, a binding table entry is recovered by querying the DHCP server or the destination host, depending on the configuration.
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-16-12/ip6f-xe-16-12-book/ip6-snooping.pdf Page 2

upvoted 4 times

Question #415 *Topic 1*

Which two components are needed for a service provider to utilize the L3VPN MPLS application? (Choose two.)

A. The P routers must be configured with RSVP.

B. The PE routers must be configured for MP-eBGP to connect to CEs.

C. The P routers must be configured for MP-iBGP toward the PE routers.

D. The PE routers must be configured for MP-iBGP with other PE routers.

E. The P and PE routers must be configured with LDP or RSVP.

Correct Answer: *DE*

HungarianDish Highly Voted 7 months, 1 week ago

Selected Answer: DE

https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/BRKMPL-1100.pdf

upvoted 5 times

inteldarvid Most Recent 4 months, 3 weeks ago

Selected Answer: DE

correct: D and E

upvoted 1 times

Question #416 *Topic 1*

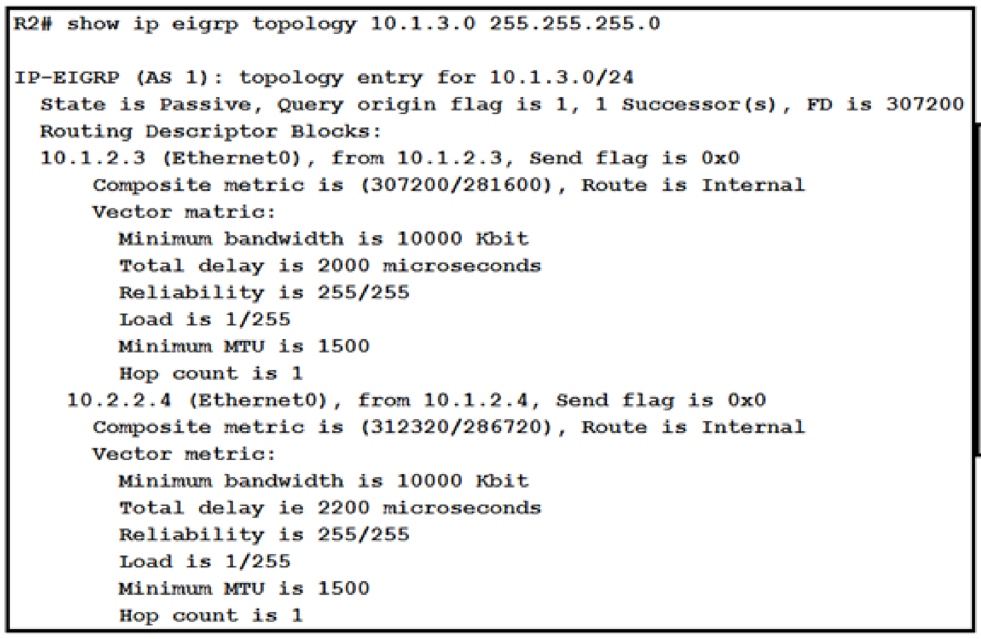

```
R2# show ip eigrp topology 10.1.3.0 255.255.255.0

IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 307200
  Routing Descriptor Blocks:
  10.1.2.3 (Ethernet0), from 10.1.2.3, Send flag is 0x0
      Composite metric is (307200/281600), Route is Internal
      Vector matric:
        Minimum bandwidth is 10000 Kbit
        Total delay is 2000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
    10.2.2.4 (Ethernet0), from 10.1.2.4, Send flag is 0x0
      Composite metric is (312320/286720), Route is Internal
      Vector metric:
        Minimum bandwidth is 10000 Kbit
        Total delay ie 2200 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
```

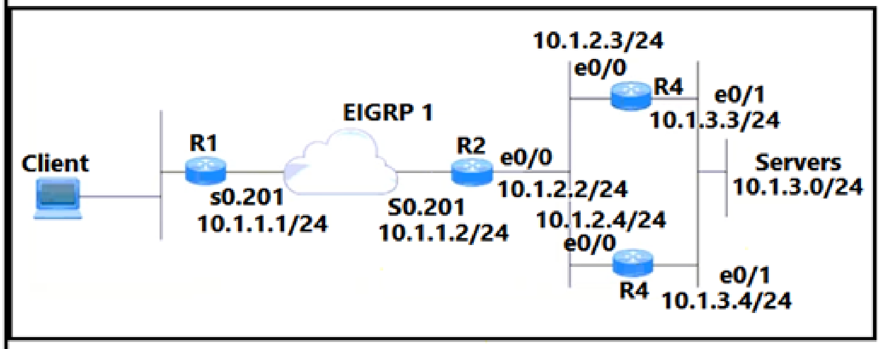

Refer to the exhibit. A network is configured for EIGRP equal-cost load balancing, but the traffic destined to the servers is not load balanced. Link metrics from router R2 to R3 and R4 are the same. Which delay value must be configured to resolve the issue?

A. 2200 on R4 E0/1

B. 200 on R3 E0/1

C. 120 on R3 E0/1

D. 120 on R4 E0/1

Correct Answer: *B*

⊟ **elmorrre** Highly Voted 9 months, 3 weeks ago

Selected Answer: C

Delay shown at a R2 is a SUM of delays R2 (e0/0) and R3 (e0/1) via R2-R3 path
and SUM of delays R2 (e0/0) and R4 (e0/1) via R2-R4 path
Thus, it is 100 (tens of microseconds) + 100 (tens of microseconds) = 200 (tens of msec) = 2000 (microseconds) via upper path, and 100 + 120 (220 tens of msec = 2200) via lower path.

For load balancing, we have to setup R3 e0/1 delay 120 (delay on an interface is setup in tens of microseconds)

Finally, we will have 2200 msec summarized delay for 10.1.3.0 subnet on R2

Labbed it in GNS3.

upvoted 13 times

⊟ **[Removed]** 8 months, 1 week ago

Have you tried in the lab to configure the exact delay as R3 on R4? In my opinion at this point the router will make the necessary calculations and because bandwidth and delay are the same it will load balance them, therefore A.

upvoted 1 times

⊟ **[Removed]** 8 months, 1 week ago

After double checking it your logic is right

upvoted 1 times

⊟ **pyrokar** Most Recent 7 months ago

Selected Answer: C

EIGRP default metric is 256*[(10^7/Min. Bandwidth)+delay/10]
Repoted distance for R4 is 286720
Min. Bandwidth in kbps is 10^4

256*[(10^7/10^4)+delay/10] = 286720
delay for R4 = 120

Hence R3 must be adjusted to the same delay -> Answer C

upvoted 1 times

⊟ **HungarianDish** 7 months, 1 week ago

Selected Answer: C

Agree with elmorre.
The scenario is taken from here, where it is explained in detail:
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/13673-14.html

upvoted 2 times

ellen_AA 11 months ago

Probably B is correct. It says, 'Link metrics from router R2 to R3 and R4 are the same' which means, the difference happen after when R2 and R3 to reach destination of 10.1.3.0/24. to make total delay equal, add 200 delay on interface facing the subnet 10.1.3.0/24 on R3.

upvoted 1 times

Patrick1234 11 months ago

I think the given answer is correct. When you configure this:

interface eth 0/1
delay 200

And do a show ip eigrp topology, the value 200 will be multiplied with 10, so it will tell you 2000, which is equal to the 2000 of the other neighbor.

upvoted 2 times

smayus 10 months ago

Why is it going to be multiplied by 10. Cumulative delay along the paths is divided by 10 at the end. This question seems to be very tricky and unclear.

upvoted 2 times

jarz 1 year ago

Selected Answer: A

it's clear the only difference between the two routers is the Total delay, R3 is 2000 and R4 is 2200.

upvoted 3 times

ellen_AA 11 months, 1 week ago

it should be 2000 on R4 E0/1 to make it equal R3 delay of 2000

upvoted 3 times

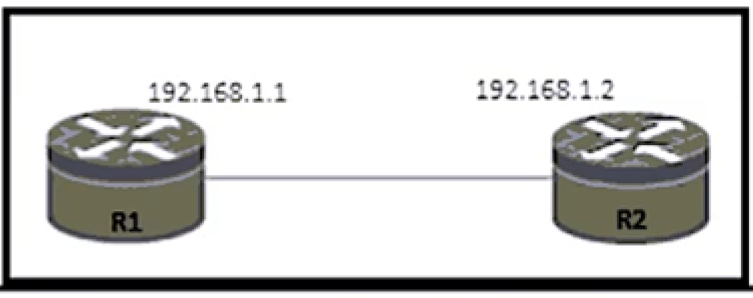

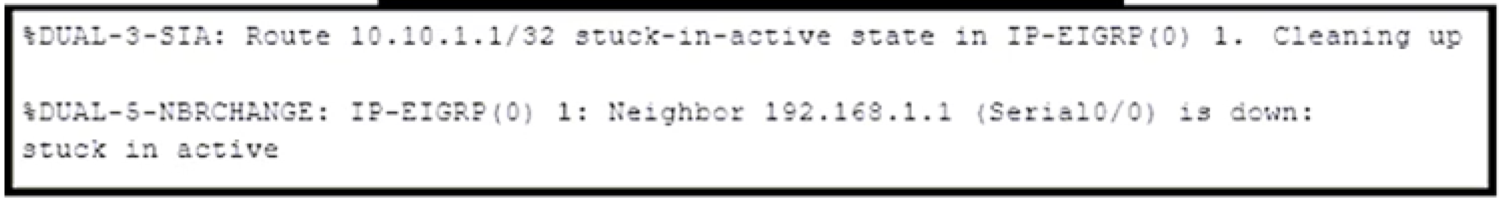

```
%DUAL-3-SIA: Route 10.10.1.1/32 stuck-in-active state in IP-EIGRP(0) 1.  Cleaning up

%DUAL-5-NBRCHANGE: IP-EIGRP(0) 1: Neighbor 192.168.1.1 (Serial0/0) is down:
stuck in active
```

Refer to the exhibit. An engineer notices a connectivity problem between routers R1 and R2. The frequency of this problem is high during peak business hours. Which action resolves the issue?

A. Increase the available bandwidth between R1 and R2.

B. Decrease the EIGRP keepalive and hold down timers on R1 and R2.

C. Increase the MTU on the interfaces that connect R1 and R2.

D. Set static EIGRP neighborship between R1 and R2.

Correct Answer: *A*

inteldarvid 4 months, 3 weeks ago

Selected Answer: A

yes, give anwser is correct

upvoted 2 times

HungarianDish 7 months, 1 week ago

Selected Answer: A

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/13676-18.html

EIGRP DUAL-3-SIA
-Missing or incorrect bandwidth interface configuration parameter
-Incorrect bandwidth configured to influence path selection

upvoted 2 times

Question #418 *Topic 1*

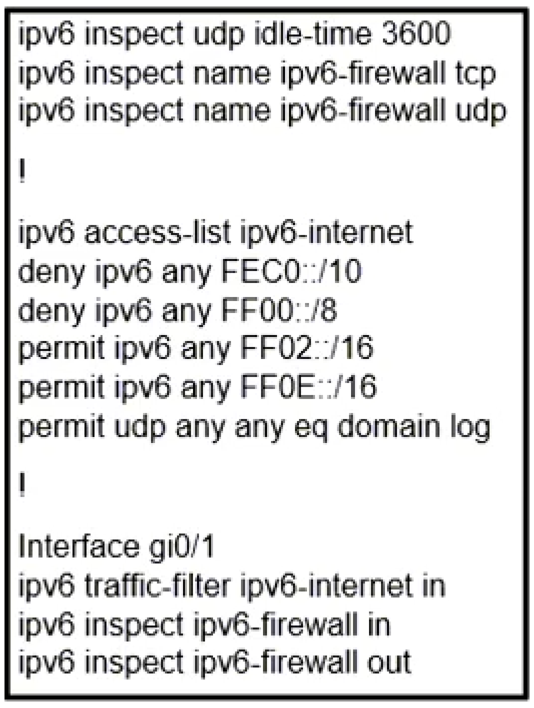

```
ipv6 inspect udp idle-time 3600
ipv6 inspect name ipv6-firewall tcp
ipv6 inspect name ipv6-firewall udp
!
ipv6 access-list ipv6-internet
deny ipv6 any FEC0::/10
deny ipv6 any FF00::/8
permit ipv6 any FF02::/16
permit ipv6 any FF0E::/16
permit udp any any eq domain log
!
Interface gi0/1
ipv6 traffic-filter ipv6-internet in
ipv6 inspect ipv6-firewall in
ipv6 inspect ipv6-firewall out
```

Refer to the exhibit. A network administrator configured name resolution for IPv6 traffic to be allowed through an inbound access list. After the access list is applied to resolve the issue, name resolution still did not work. Which action does the network administrator take to resolve the name resolution problem?

A. Add permit udp any eq domain 53 any log in the access list

B. Remove ipv6 inspect ipv6-firewall in from interface gi0/1

C. Add permit udp any eq domain any log in the access list

D. Inspect ipv6 inspect name ipv6-firewall udp 53 in global config

Correct Answer: *C*

DeWalt95 1 week ago

C makes the most sense - will permit replies from the DHCP server

upvoted 1 times

slcc99 5 months ago

Yes,C.To be exact, TCP should also be allowed.

https://www.cisco.com/c/en/us/support/docs/ip/access-lists/26448-ACLsamples.html#anc18:~:text=This%20configuration%20permits%20TCP%20traffic%20with%20destination%20port%20value%2053.%20The%20implicit%20deny%20all%20clause%20at%20the%20end%20of%20an%20ACL%20denies%20all%20other%20traffic%2C%20which%20does%20not%20match%20the%20permit%20clauses.&text=access%2Dlist%20102%20permit%20udp%20any%20eq%20domain%20any

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

yes correct C

upvoted 1 times

Doggolover15 7 months ago

https://community.cisco.com/t5/ipv6/ipv6-acl-problem-with-dns/td-p/2271733

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: C

Please, see explanation here:
https://networkengineering.stackexchange.com/questions/46248/configuring-acl-for-dns

ACE - Client to server UDP case for DNS queries:
permit udp any any eq domain

Reply ACE for DNS replies with the A-record:
permit udp any eq domain any

Scenario is taken from:
https://community.cisco.com/t5/ipv6/ipv6-acl-problem-with-dns/td-p/2271733

upvoted 3 times

HungarianDish 7 months, 1 week ago

Not "B". "ipv6 inspect ipv6-firewall in" on the interface facing the internet was unnecessary, however it did not cause the specific issue in the question. Please, see:
https://community.cisco.com/t5/ipv6/ipv6-acl-problem-with-dns/td-p/2271733

upvoted 2 times

azzawim 9 months ago

Selected Answer: B

answer is B

upvoted 1 times

Typovy 9 months, 1 week ago

Can anyone explain?

upvoted 1 times

Question #419 *Topic 1*

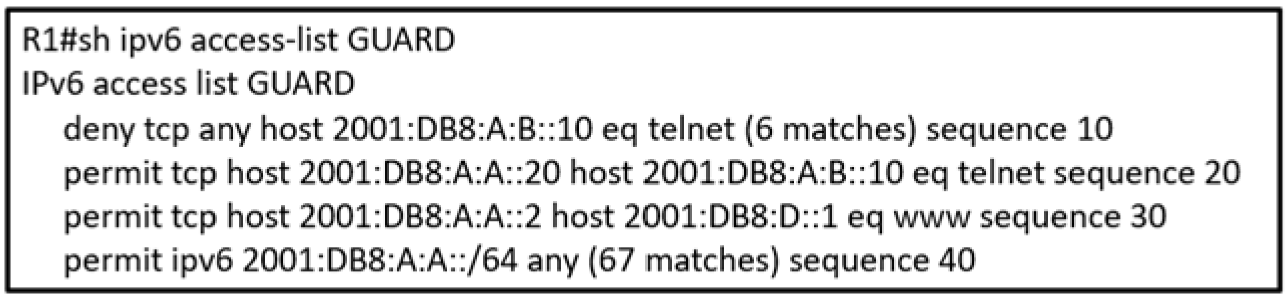

```
R1#sh ipv6 access-list GUARD
IPv6 access list GUARD
    deny tcp any host 2001:DB8:A:B::10 eq telnet (6 matches) sequence 10
    permit tcp host 2001:DB8:A:A::20 host 2001:DB8:A:B::10 eq telnet sequence 20
    permit tcp host 2001:DB8:A:A::2 host 2001:DB8:D::1 eq www sequence 30
    permit ipv6 2001:DB8:A:A::/64 any (67 matches) sequence 40
```

Refer to the exhibit. PC2 is directly connected to R1. A user at PC2 cannot Telnet to 2001:db8:a:b::10. The user can ping 2001:db8:a:b::10 and receive DHCP-related information from the DHCP server.

Which action resolves the issue?

A. Remove sequence 30 and put it back as sequence 5.

B. Remove sequence 10 and put it back as sequence 25.

C. Remove sequence 40 and put it back as sequence 15.

D. Remove sequence 20 and put it back as sequence 45.

Correct Answer: *B*

inteldarvid 5 months, 1 week ago

Selected Answer: B

B correct

upvoted 1 times

HungarianDish 7 months, 1 week ago

I assume, 2001:db8:a:a::20 is PC2.

upvoted 2 times

Question #420

Topic 1

What is an advantage of implementing BFD?

A. BFD is deployed without the need to run any routing protocol.

B. BFD provides faster updates for any flapping route.

C. BFD provides millisecond failure detection.

D. BFD provides better capabilities to maintain the routing table.

Correct Answer: *C*

inteldarvid 4 months, 3 weeks ago

Selected Answer: C

the give answer is correct

upvoted 1 times

Question #421

Topic 1

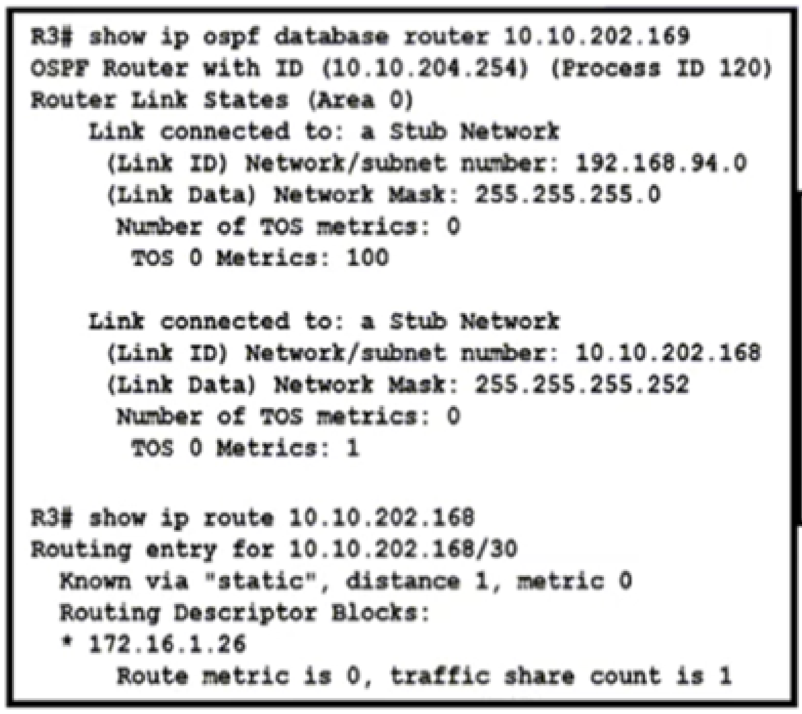

```
R3# show ip ospf database router 10.10.202.169
OSPF Router with ID (10.10.204.254) (Process ID 120)
Router Link States (Area 0)
    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 192.168.94.0
     (Link Data) Network Mask: 255.255.255.0
      Number of TOS metrics: 0
       TOS 0 Metrics: 100

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 10.10.202.168
     (Link Data) Network Mask: 255.255.255.252
      Number of TOS metrics: 0
       TOS 0 Metrics: 1

R3# show ip route 10.10.202.168
Routing entry for 10.10.202.168/30
  Known via "static", distance 1, metric 0
  Routing Descriptor Blocks:
  * 172.16.1.26
      Route metric is 0, traffic share count is 1
```

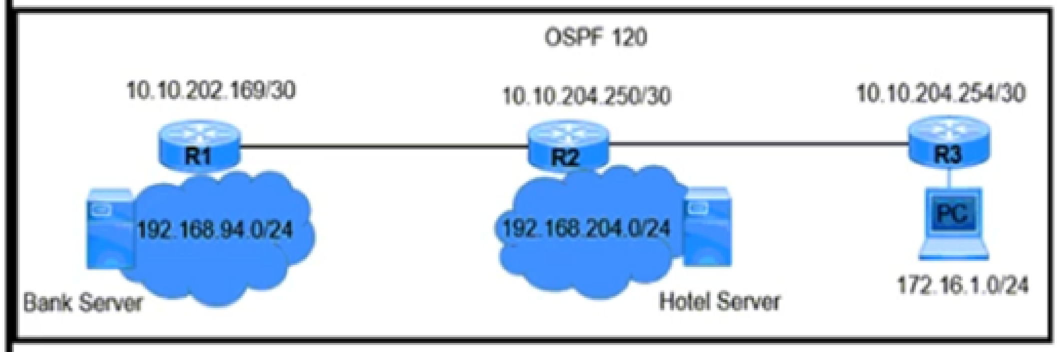

Refer to the exhibit. A network engineer finds that PC1 is accessing the hotel website to do the booking but fails to make payment. Which action resolves the issue?

A. Increase the AD to 200 of static route 192.168.94.0 on R3.

B. Configure a reverse route on R1 for PC1 172.16.1.0/24.

C. Decrease the AD to 5 of OSPF route 192.168.94.0 on R1.

D. Allow stub network 10.10.202.168/30 on router R3 OSPF.

Correct Answer: *B*

ellen_AA Highly Voted 11 months, 1 week ago

Since the Bank routes are in stub area (which blocks type 4 & 5 LSA), it won't be able to receive R3 redistributed static route into OSPF, thus R1 is not able to communicate back to R3. Two options are envisageable:
- configure a reverse static route on R1 back to R3
- configure the Bank router area as Regular or NSSA area which allow redistributed routes to come in.

upvoted 7 times

JoeyT 6 months ago

ur explaination might be right, but first of all, why R3 has a static route to 172.16.1.26 for destination IPs

upvoted 4 times

AlexInShort12 Most Recent 5 days ago

Selected Answer: A

I would say A, based on the topology that all three router are in the same Stub area 120.
Increasing the static route AD would make sure to use the route that is display in DB router command which is good.

upvoted 1 times

Malasxd 7 months ago

Selected Answer: B

B is right

upvoted 2 times

Question #422 *Topic 1*

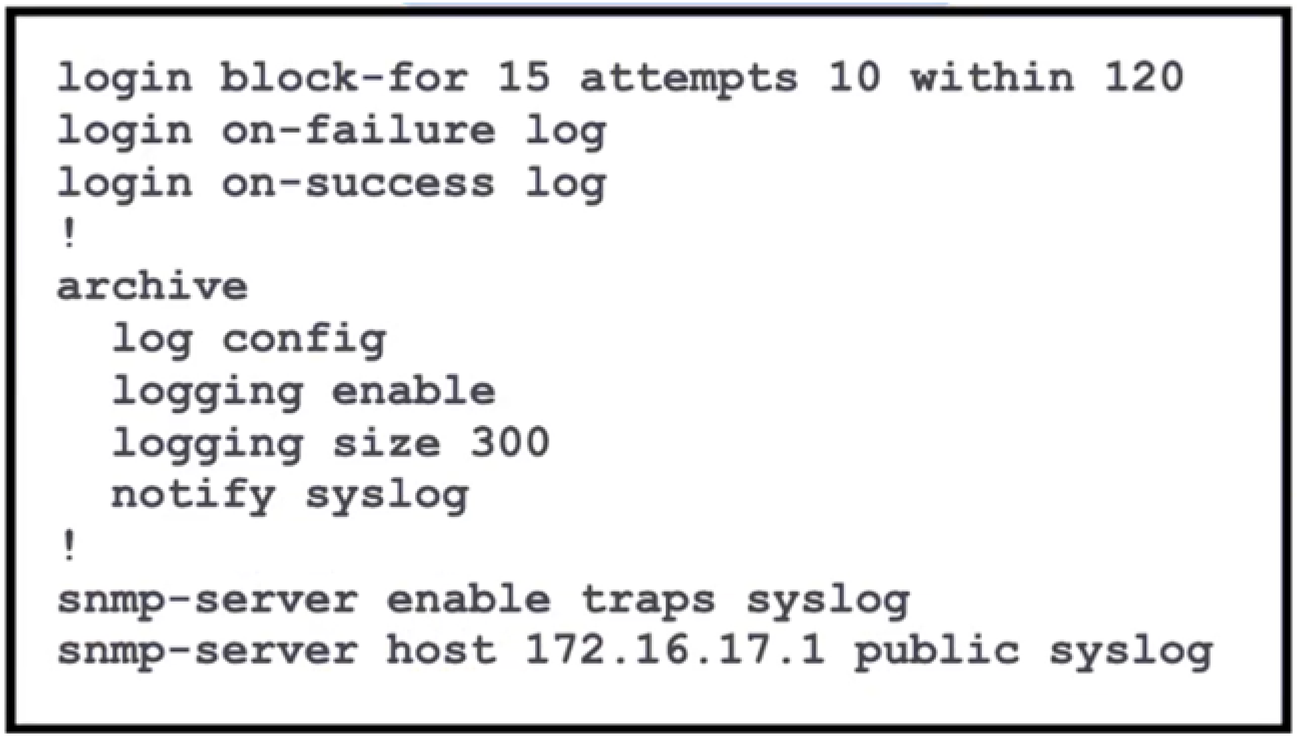

```
login block-for 15 attempts 10 within 120
login on-failure log
login on-success log
!
archive
  log config
  logging enable
  logging size 300
  notify syslog
!
snmp-server enable traps syslog
snmp-server host 172.16.17.1 public syslog
```

Refer to the exhibit. The administrator can see the traps for the failed login attempts, but cannot see the traps of successful login attempts. Which action fixes this issue?

A. Configure logging history 4.

B. Configure logging history 3.

C. Configure logging history 2.

D. Configure logging history 5.

Correct Answer: *D*

ellen_AA 11 months, 1 week ago

#logging history 4, is the default behavior, and doesn't show the intended result.
#logging history 5, is the LOG_NOTICE

upvoted 3 times

Question #423 *Topic 1*

An engineer creates a default static route on a router with a next hop of 10.1.1.1. On inspection, the engineer finds the router has two VRFs, Red and Blue. The next hop is valid for both VRFs and exists in each assigned VRF. Which configuration achieves connectivity?

A. ip route vrf Red 0.0.0.0 0.0.0.0 10.1.1.1
ip route vrf Blue 0.0.0.0 0.0.0.0 10.1.1.1

B. ip route vrf BLUE 0.0.0.0 255.255.255.255 10.1.1.1
ip route vrf RED 0.0.0.0 255.255.255.255 10.1.1.1

C. ip route vrf Red 0.0.0.0 255.255.255.255 10.1.1.1

D. ip route vrf Blue 0.0.0.0 255.255

Correct Answer: *A*

inteldarvid 4 months, 3 weeks ago

Selected Answer: A

the give anwser is correct

upvoted 1 times

ellen_AA 11 months, 1 week ago

A is correct!

upvoted 3 times

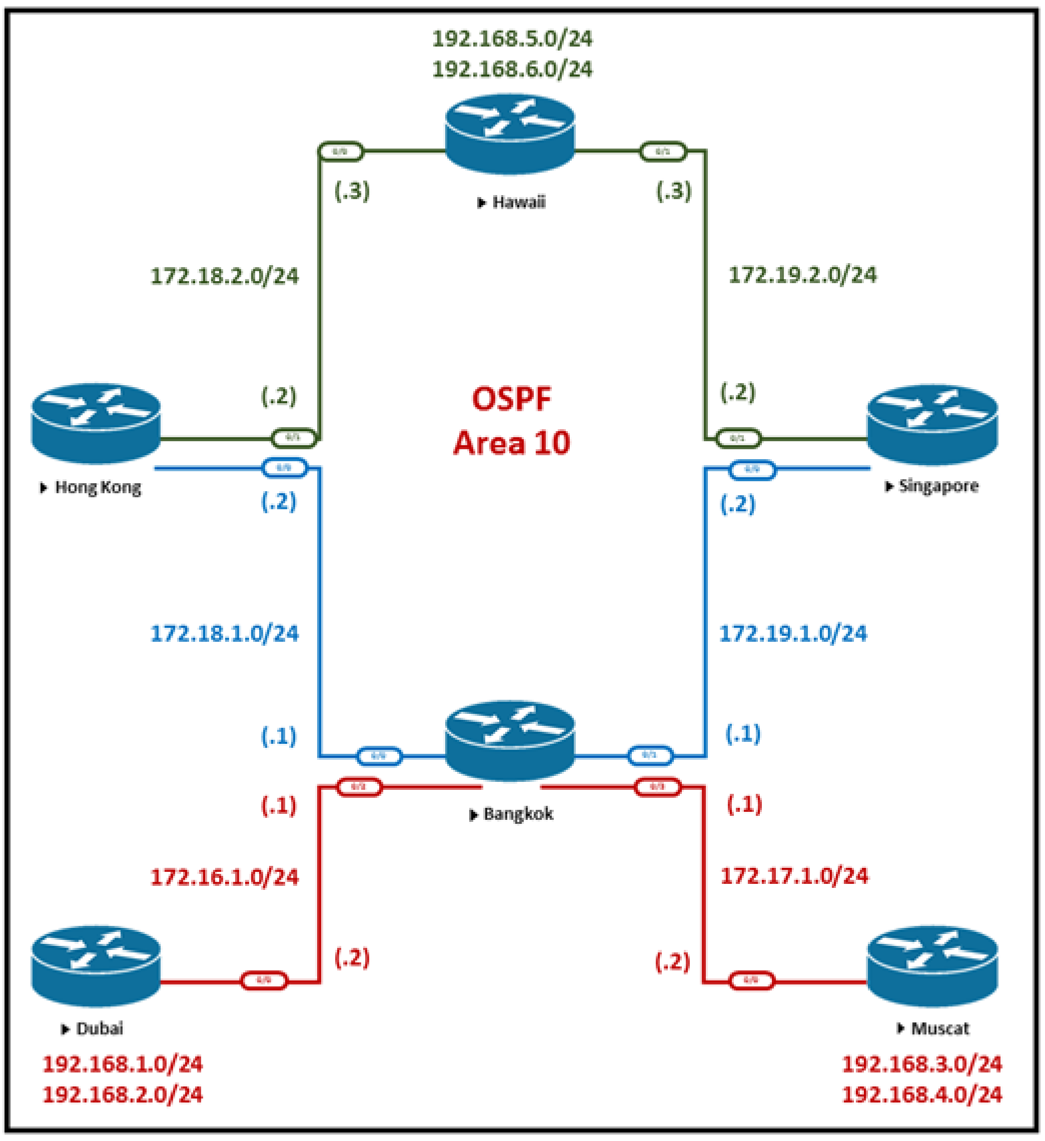

Refer to the exhibit. Bangkok is using ECMP to reach to the 192.168.5.0/24 network. The administrator must configure Bangkok in such a way that Telnet traffic from 192.168.3.0/24 and 192.168.4.0/24 networks use the Hong Kong router as the preferred route. Which set of configurations accomplishes this task?

A. access-list 101 permit tcp 192.168.3.0 0.0.0.255 192.168.5.0 0.0.0.255 eq 23 access-list 101 permit tcp 192.168.4.0 0.0.0.255 192.168.5.0 0.0.0.255 eq 23
!
route-map PBR1 permit 10
match ip address 101
set ip next-hop 172.18.1.2
!
interface Ethernet0/1
ip policy route-map PBR1

B. access-list 101 permit tcp 192.168.3.0 0.0.0.255 192.168.5.0 0.0.0.255 eq 23 access-list 101 permit tcp 192.168.4.0 0.0.0.255 192.168.5.0 0.0.0.255 eq 23
!
route-map PBR1 permit 10
match ip address 101
set ip next-hop 172.18.1.2
!
interface Ethernet0/3
ip policy route-map PBR1

C. access-list 101 permit ip 192.168.3.0 0.0.0.255 192.168.5.0 0.0.0.255 access-list 101 permit ip 192.168.4.0 0.0.0.255 192.168.5.0 0.0.0.255
!
route-map PBR1 permit 10
match ip address 101
set ip next-hop 172.18.1.2
!
interface Ethernet0/3
ip policy route-map PBR1

D.

```
access-list 101 permit ip 192.168.3.0 0.0.0.255 192.168.5.0 0.0.0.255 access-list 101 permit ip 192.168.4.0 0.0.0.255 192.168.5.0 0.0.0.255
!
route-map PBR1 permit 10
match ip address 101
set ip next-hop 172.18.1.2
!
interface Ethernet0/1
ip policy route-map PBR1
```

Correct Answer: *B*

Hummer1 5 months ago

Bangkok router interfaces are:

E0/0 E0/1
Bangkok
E0/2 E0/3

upvoted 2 times

inteldarvid 5 months, 1 week ago

i dont see interface

upvoted 4 times

HungarianDish 7 months, 1 week ago

Selected Answer: B

https://networklessons.com/cisco/ccie-routing-switching/how-to-configure-policy-based-routing

upvoted 1 times

Question #425

Topic 1

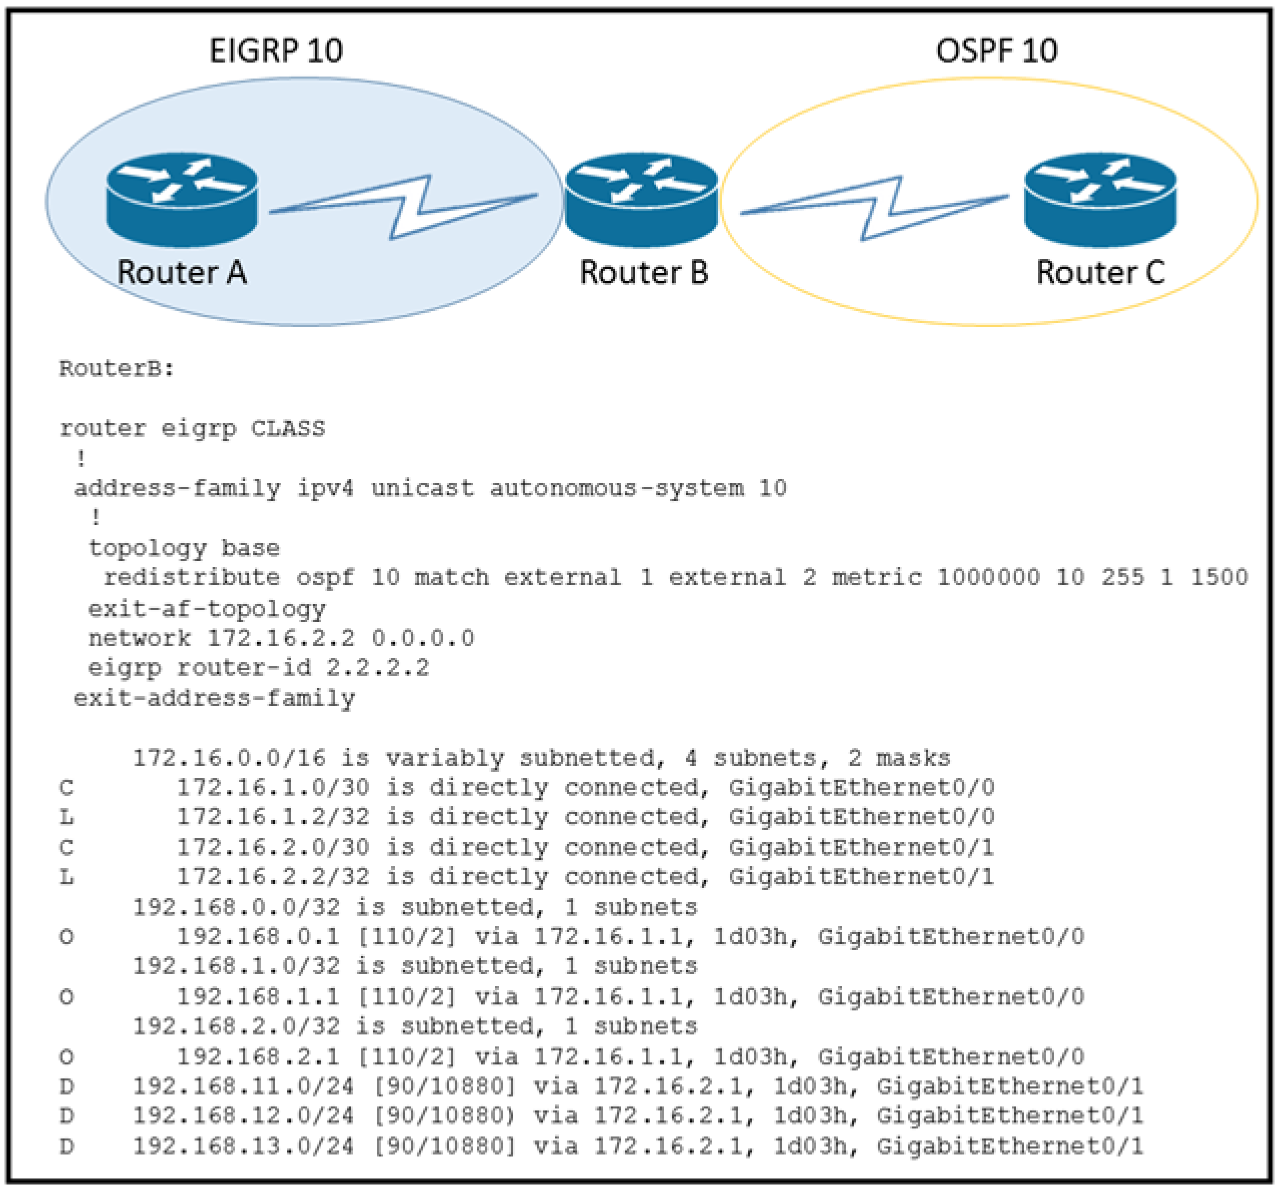

```
RouterB:

router eigrp CLASS
 !
 address-family ipv4 unicast autonomous-system 10
  !
  topology base
   redistribute ospf 10 match external 1 external 2 metric 1000000 10 255 1 1500
  exit-af-topology
  network 172.16.2.2 0.0.0.0
  eigrp router-id 2.2.2.2
 exit-address-family

     172.16.0.0/16 is variably subnetted, 4 subnets, 2 masks
C       172.16.1.0/30 is directly connected, GigabitEthernet0/0
L       172.16.1.2/32 is directly connected, GigabitEthernet0/0
C       172.16.2.0/30 is directly connected, GigabitEthernet0/1
L       172.16.2.2/32 is directly connected, GigabitEthernet0/1
     192.168.0.0/32 is subnetted, 1 subnets
O       192.168.0.1 [110/2] via 172.16.1.1, 1d03h, GigabitEthernet0/0
     192.168.1.0/32 is subnetted, 1 subnets
O       192.168.1.1 [110/2] via 172.16.1.1, 1d03h, GigabitEthernet0/0
     192.168.2.0/32 is subnetted, 1 subnets
O       192.168.2.1 [110/2] via 172.16.1.1, 1d03h, GigabitEthernet0/0
D    192.168.11.0/24 [90/10880] via 172.16.2.1, 1d03h, GigabitEthernet0/1
D    192.168.12.0/24 [90/10880) via 172.16.2.1, 1d03h, GigabitEthernet0/1
D    192.168.13.0/24 [90/10880] via 172.16.2.1, 1d03h, GigabitEthernet0/1
```

Refer to the exhibit. An engineer configured route exchange between two different companies for a migration project. EIGRP routes were learned in router C, but no OSPF routes were learned in router A. Which configuration allows router A to receive OSPF routes?

A. (config-router-af-topology)#no redistribute ospf 10 match external 1 external 2 metric 1000000 10 255 1 1500

B. (config-router-af)#redistribute ospf 10 1000000 10 255 1 1500

C. (config-router-af-topology)#redistribute connected

D. (config-router-af-topology)#redistribute ospf 10 metric 1000000 10 255 1 1500

Correct Answer: *D*

⊟ **ellen_AA** Highly Voted 11 months, 1 week ago

Engineer redistribute external 1 external 2, which don't exist by looking into the exhibit.
Answer is D, because it uses the default behavior of redistribution of OSPF routes into EIGRP. Only intra-area 'O' and inter-area '0 IA' are redistributed using this command!

upvoted 7 times

⊟ **inteldarvid** Most Recent 5 months, 1 week ago

Selected Answer: D

the answer correct:

```
IOU6(config-router-af-topology)#redistribute ospf 10 ?
match Redistribution of OSPF routes
metric Metric for redistributed routes
route-map Route map reference
<cr>

IOU6(config-router-af-topology)#redistribute ospf 10 metr
IOU6(config-router-af-topology)#redistribute ospf 10 metric ?
<1-4294967295> Bandwidth metric in Kbits per second

IOU6(config-router-af-topology)#redistribute ospf 10 metric
```

upvoted 1 times

Question #426

Topic 1

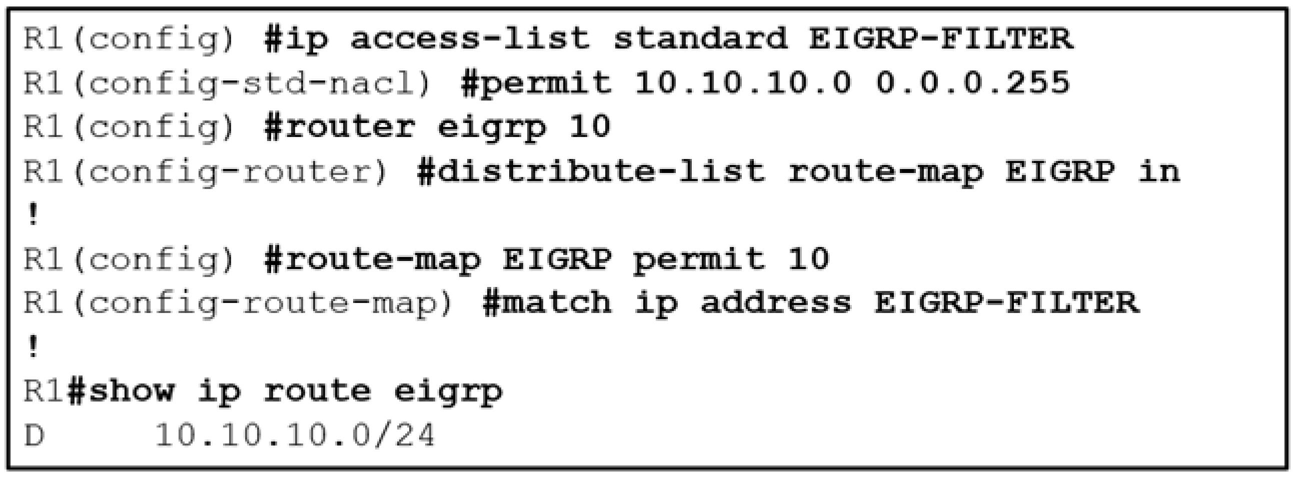

```
R1(config) #ip access-list standard EIGRP-FILTER
R1(config-std-nacl) #permit 10.10.10.0 0.0.0.255
R1(config) #router eigrp 10
R1(config-router) #distribute-list route-map EIGRP in
!
R1(config) #route-map EIGRP permit 10
R1(config-route-map) #match ip address EIGRP-FILTER
!
R1#show ip route eigrp
D     10.10.10.0/24
```

Refer to the exhibit. An engineer must filter incoming EIGRP updates to allow only a set of specific prefixes. The distribute list is tested, and it filters out all routes except network 10.10.10.0/24. How should the engineer temporarily allow all prefixes to be learned by the router again without adjusting the existing access list?

A. A permit any statement should be added before completing the ACL with the required prefixes, and then the permit any statement can be removed.

B. A permit 20 statement should be added before completing the ACL with the required prefixes, and then the permit 20 statement can be removed.

C. A continue statement should be added within the permit 10 statement before completing the ACL with the required prefixes, and then the continue statement can be removed.

D. An extended access list must be used instead of a standard access list to accomplish the task.

Correct Answer: *B*

⊟ **[Removed]** 4 months, 3 weeks ago

Selected Answer: B

Requirement: temporarily allow all routes without modifying the ACCESS-LIST
D. is wrong, this specifies directly modifying access list
A. is wrong, there is no such command as "permit any" on a route-map
C. is wrong, the continue statement under route-map, as I understand it, means that it should continue to process the route-map entry for the next statement in it.

B. is correct, we need to override the implicit deny at the end of a route-map list with a statement that permits all routes. This will allow the engineer to modify the ACL without causing a longer outage.

upvoted 1 times

⊟ **Brand** 3 months, 3 weeks ago

My friend, both A and B is talking about adding permit any to ACL.

upvoted 2 times

⊟ **Gedson** 5 months ago

Selected Answer: A

When we first read the requirement "... without adjusting the existing access list", we think the best solution is to add the statement "route-map EIGRP permit 20" (without any "match" statement) at the end of the route-map but there is no such choice.

upvoted 4 times

⊟ **HungarianDish** 7 months, 1 week ago

Selected Answer: B

Tested the scenario in CML.

upvoted 1 times

⊟ **HungarianDish** 7 months, 1 week ago

It is about adding and removing this part of the route-map:
route-map EIGRP permit 20

upvoted 1 times

⊟ **HungarianDish** 7 months, 1 week ago

"permit statement that when used without any match statements, means “match all” effectively permitting all other routes"
http://notthenetwork.me/blog/2013/07/24/ccnp-route-study-eigrp-route-filtering/

upvoted 1 times

GReddy2323 10 months, 2 weeks ago

Can someone please explain? It says without adjusting the existing access list... how can it be B?

upvoted 1 times

Typovy 9 months, 1 week ago

This mean permit statement in routemap

upvoted 1 times

Question #427 *Topic 1*

What is a function of the IPv6 DHCP Guard feature for DHCP messages?

A. If the device is configured as a DHCP server, no message is switched.

B. All client messages are always switched regardless of the device role.

C. It blocks only DHCP request messages.

D. Only access lists are supported for matching traffic.

Correct Answer: *B*

DUBC89x Highly Voted 👍 1 year ago

Selected Answer: B

DHCPv6 Guard Overview

The DHCPv6 Guard feature blocks reply and advertisement messages that come from unauthorized DHCP servers and relay agents. Packets are classified into one of the three DHCP type messages. All client messages are always switched regardless of device role. DHCP server messages are only processed further if the device role is set to server. Further processing of server messages includes DHCP server advertisements(for source validation and server preference) and DHCP server replies (for permitted prefixes). If the device is configured as a DHCP server, all the messages need to be switched, regardless of the device role configuration.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_dhcp/configuration/15-sy/dhcp-15-sy-book/ip6-dhcpv6-guard.pdf

upvoted 6 times

Hummer1 Most Recent 🕒 5 months ago

Selected Answer: B

Correct

upvoted 1 times

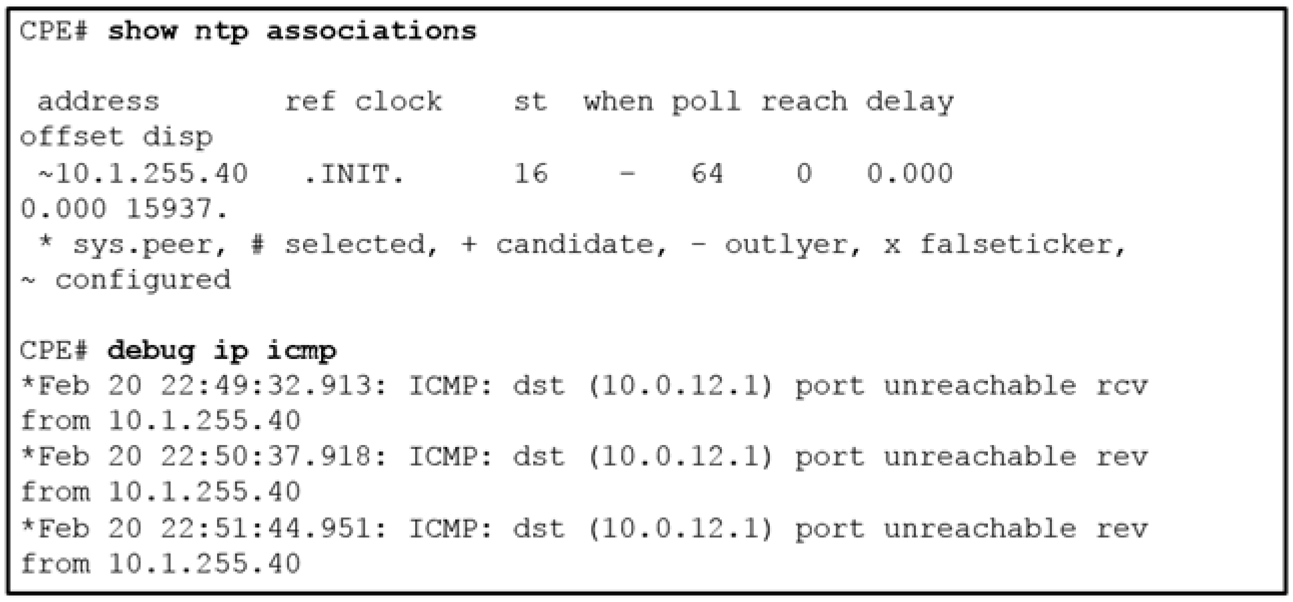

```
CPE# show ntp associations

 address        ref clock    st  when poll reach delay
offset disp
 ~10.1.255.40   .INIT.       16    -   64    0   0.000
0.000 15937.
 * sys.peer, # selected, + candidate, - outlyer, x falseticker,
~ configured

CPE# debug ip icmp
*Feb 20 22:49:32.913: ICMP: dst (10.0.12.1) port unreachable rcv
from 10.1.255.40
*Feb 20 22:50:37.918: ICMP: dst (10.0.12.1) port unreachable rev
from 10.1.255.40
*Feb 20 22:51:44.951: ICMP: dst (10.0.12.1) port unreachable rev
from 10.1.255.40
```

Refer to the exhibit. An administrator is troubleshooting a time synchronization problem for the router's time to another Cisco IOS XE-based device that has recently undergone security hardening.

Which action resolves the issue?

A. NTP service is disabled and must be enabled on 10.1.255.40.

B. Ensure that the CPE router has a valid route to 10.1.255.40 for NTP and rectify if not reachable.

C. Allow NTP in the ingress ACL on 10.1.255.40 by permitting UDP destined to port 123.

D. Allow NTP in the ingress ACL on 10.1.255.40 by permitting TCP destined to port 123.

Correct Answer: *A*

aqwsdfghjklp 2 weeks, 3 days ago

https://kb.meinbergglobal.com/kb/time_sync/ntp/ntp_debugging/ntp_debugging_unreachable_time_sources

Why not B?

upvoted 1 times

Patrick1234 10 months, 4 weeks ago

Selected Answer: A

Sorry, i was wrong, also labbed it and Zizu007 is right.

When filtered by access list:

Jan 30 10:22:13.908: ICMP: dst (181.16.2.6) administratively prohibited unreachable rcv from 181.16.2.5

When NTP is turned off on master:

*Jan 30 10:13:20.287: ICMP: dst (181.16.2.6) port unreachable rcv from 181.16.2.5

So in this case NTP needs to be enabled.

upvoted 4 times

Patrick1234 11 months ago

Correct answer is C. Key word in the question is "security hardening".

upvoted 1 times

Zizu007 11 months, 2 weeks ago

Selected Answer: A

Correct!

if filtered by ACL this msg will show up:
R3#
ICMP: dst (10.0.12.1) administratively prohibited unreachable rcv from 10.1.255.40
R3#

upvoted 3 times

juliop 12 months ago

Why A? for me is C the correct Answer

upvoted 1 times

Zizu007 11 months, 2 weeks ago

if filtered by ACL this msg will show up:

```
R3#
ICMP: dst (10.0.12.1) administratively prohibited unreachable rcv from 10.1.255.40
R3#
```

upvoted 2 times

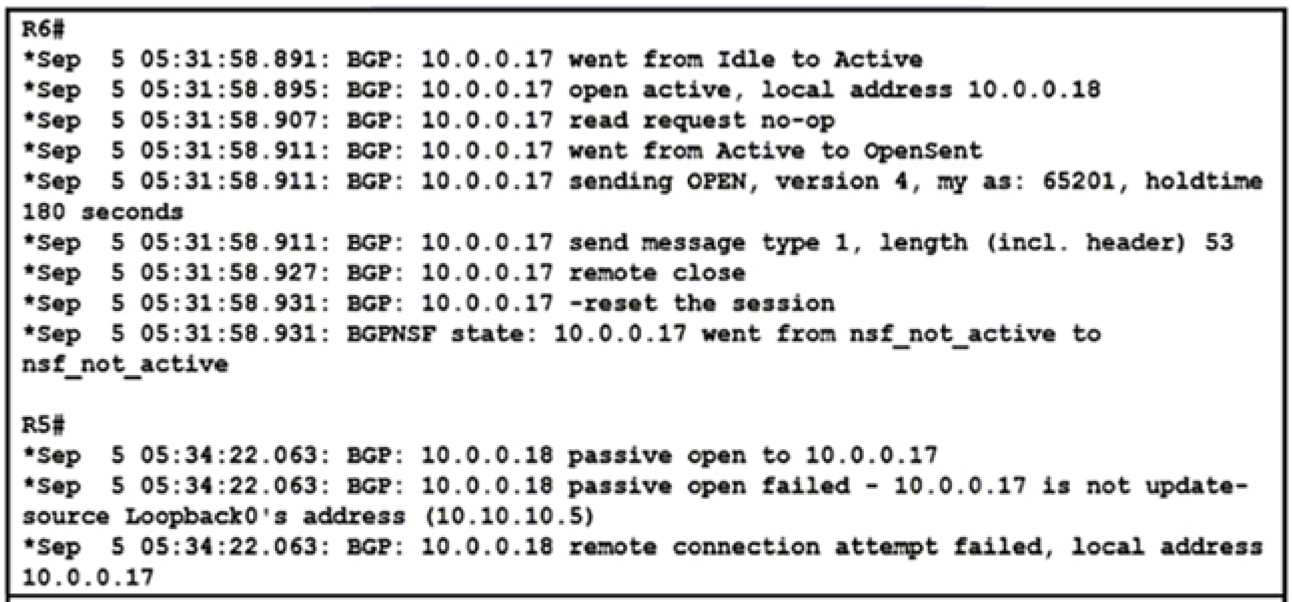

```
R6#
*Sep  5 05:31:58.891: BGP: 10.0.0.17 went from Idle to Active
*Sep  5 05:31:58.895: BGP: 10.0.0.17 open active, local address 10.0.0.18
*Sep  5 05:31:58.907: BGP: 10.0.0.17 read request no-op
*Sep  5 05:31:58.911: BGP: 10.0.0.17 went from Active to OpenSent
*Sep  5 05:31:58.911: BGP: 10.0.0.17 sending OPEN, version 4, my as: 65201, holdtime
180 seconds
*Sep  5 05:31:58.911: BGP: 10.0.0.17 send message type 1, length (incl. header) 53
*Sep  5 05:31:58.927: BGP: 10.0.0.17 remote close
*Sep  5 05:31:58.931: BGP: 10.0.0.17 -reset the session
*Sep  5 05:31:58.931: BGPNSF state: 10.0.0.17 went from nsf_not_active to
nsf_not_active

R5#
*Sep  5 05:34:22.063: BGP: 10.0.0.18 passive open to 10.0.0.17
*Sep  5 05:34:22.063: BGP: 10.0.0.18 passive open failed - 10.0.0.17 is not update-
source Loopback0's address (10.10.10.5)
*Sep  5 05:34:22.063: BGP: 10.0.0.18 remote connection attempt failed, local address
10.0.0.17
```

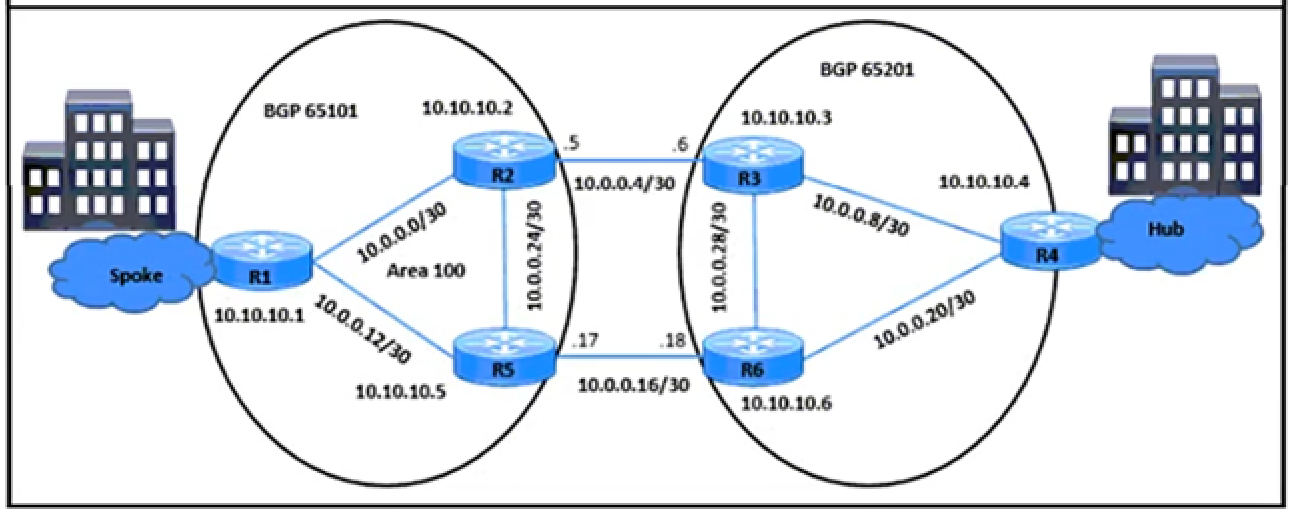

Refer to the exhibit. The traffic from spoke to hub is dropping. The operations team observes:

- R2-R3 link is down due to the fiber cut.
- R2 and R5 receive traffic from R1 in AS 65101.
- R3 and R6 receive traffic from R4 in AS 65201.

Which configuration resolves the issue?

A. R5(config)#router bgp 65101 -
R5(config-router)#neighbor 10.10.10.6 remote-as 65201
R5(config-router)#neighbor 10.10.10.6 update-source Loopback0
R5(config-router)#neighbor 10.10.10.6 ebgp-multihop 3

B. R5(config)#router bgp 65101 -
R5(config-router)#no neighbor 10.0.0.18 update-source Loopback0

C. R6(config)#router bgp 65101 -
R6(config-router)#no neighbor 10.0.0.17 update-source Loopback0

D. R6(config)#router bgp 65201 -
R6(config-router)#neighbor 10.10.10.5 remote-as 65101
R6(config-router)#neighbor 10.10.10.5 update-source Loopback0
R6(config-router)#neighbor 10.10.10.5 ebgp-multihop 3

Correct Answer: *B*

HungarianDish 7 months, 1 week ago

Selected Answer: B

Reproduced error in CML lab with a simple BGP config:
R5
router bgp 65101
neighbor 10.0.0.18 remote-as 65201
network 10.10.10.5 mask 255.255.255.255
!debug ip bgp -> enabled on both routers for viewing the error
!When adding this incorrect line the same error comes as in the output:
neighbor 10.0.0.18 update-source Loopback0
R6

```
router bgp 65201
neighbor 10.0.0.17 remote-as 65101
network 10.10.10.6 mask 255.255.255.255
```

upvoted 2 times

HungarianDish 7 months, 1 week ago

Config for update-source loopback0:

```
R5
router bgp 65101
neighbor 10.10.10.6 remote-as 65201
neighbor 10.10.10.6 update-source Loopback0
neighbor 10.10.10.6 ebgp-multihop 2
ip route 10.10.10.6 255.255.255.255 10.0.0.18
!Static Route to reach Loopback0 of R6
R6
router bgp 65201
neighbor 10.10.10.5 remote-as 65101
neighbor 10.10.10.5 update-source Loopback0
neighbor 10.10.10.5 ebgp-multihop 2
ip route 10.10.10.5 255.255.255.255 10.0.0.17
!Static Route to reach Loopback0 of R5
```

upvoted 1 times

Titini 10 months ago

Selected Answer: B

The error message "BGP: 10.0.0.18 passive open failed - 10.0.0.17 is not update-source Loopback0's address (10.10.10.5)" indicates that there was a problem with establishing a BGP session between two routers with IP addresses 10.0.0.18 and 10.0.0.17.
The message specifically states that the passive open attempt from the device with IP address 10.0.0.18 has failed because the device with IP address 10.0.0.17 is not using the correct source IP address when sending updates. Instead of using the IP address of the Loopback0 interface, which is 10.10.10.5, it is using a different IP address.
In this case, it appears that the device with IP address 10.0.0.17 is using a different source IP address than the one expected by the device with IP address 10.0.0.18.
Specifically, the device with IP address 10.0.0.17 should be configured to use the IP address of its Loopback0 interface (10.10.10.5 in this case) as the source address for BGP updates when communicating with the device at IP address 10.0.0.18.

upvoted 3 times

Lilienen 10 months, 1 week ago

Selected Answer: B

Could someone actually explain this question please? I feel like we don't have enough information from the question itself.
For me, B makes sense. Current situation is, that R5 is using 10.10.10.5 address for BGP communication with R6. Also, R6 is using 10.0.0.18 for BGP communication with R5 and it's trying to open a session with R5's address 10.0.0.17. Which means:

A - won't help, because then R6 will not respond as it's not using its Lo0 address for communication
B - will help, because we remove Lo0 as update source from R5, so R6 can use R5's address 10.0.0.17 and R5 will reply
C - won't help, because R6 is clearly not using Lo0 for communication, so this command is redundant
D - won't help, because we would make R6 use R5's Lo0 address for communication, but we would also change R6's communication address too, so then there would be a mismatch with R5's neighbor config

upvoted 2 times

pitcholo 10 months, 2 weeks ago

B IS CORRECT

upvoted 1 times

ellen_AA 11 months, 1 week ago

I think it's A, it says in the exhibit of R5, 10.0.0.17 is not update source loopback0. We need to to update source loopback and update ebgp multihop to 3.

upvoted 3 times

Zizu007 11 months, 2 weeks ago

Selected Answer: B

Correct!

upvoted 2 times

Question #430 *Topic 1*

What is a function of an end device configured with DHCPv6 guard?

A. If it is configured as a client, messages are switched regardless of the assigned role.

B. If it is configured as a client, only DHCP requests are permitted.

C. If it is configured as a relay agent, only prefix assignments are permitted.

D. If it is configured as a server, only prefix assignments are permitted.

Correct Answer: *A*

⊟ **ZamanR** 4 days, 21 hours ago

A answer is correct
Explanation

The DHCPv6 Guard feature blocks reply and advertisement messages that come from unauthorized

DHCP servers and relay agents.

Packets are classified into one of the three DHCP type messages. All client messages are always

switched regardless of device role. DHCP server messages are only processed further if the device role

is set to server. Further processing of server messages includes DHCP server advertisements (for

source validation and server preference) and DHCP server replies (for permitted prefixes).

If the device is configured as a DHCP server, all the messages need to be switched, regardless of the

device role configuration.

upvoted 1 times

⊟ **Fenix7** 3 months, 1 week ago

Correct option is A.

The DHCPv6 Guard classifies the information into one of the three DHCP type messages (client message, server message, and relay message), and takes action depending on the device role. All client messages are switched regardless of the device role, and the DHCP server messages are only processed further if the device role is set to server.

upvoted 1 times

⊟ **chris110** 3 months, 2 weeks ago

Selected Answer: B

B. If it is configured as a client, only DHCP requests are permitted.

When an end device is configured with DHCPv6 guard, its role is to act as a DHCPv6 client. DHCPv6 guard is a security feature in IPv6 networks that helps prevent rogue DHCPv6 servers from providing unauthorized IPv6 configuration information to clients.

With DHCPv6 guard enabled on an end device configured as a client, it will only allow DHCPv6 requests to be sent and received. This means that the device will ignore any unauthorized DHCPv6 server responses or advertisements, helping to ensure that IPv6 configuration information is obtained only from trusted and authorized DHCPv6 servers on the network. This is important for maintaining network security and preventing potential misconfigurations or security risks.

upvoted 2 times

⊟ **MNem** 3 months, 2 weeks ago

Selected Answer: B

Correct answer is B.

If it is configured as a client, only requests are permitted. The link HungarianDish linked states that all client messages are forwarded, not that all packets are forwarded when a port is configured as a client.

All packets are forwarded when the port is configured a server. "If the device is configured as a DHCP server, all the messages need to be switched, regardless of the device role configuration."

upvoted 2 times

⊟ **inteldarvid** 4 months, 3 weeks ago

Selected Answer: A

correct option A

upvoted 2 times

⊟ **HungarianDish** 7 months, 1 week ago

Selected Answer: A

Q#427 - DUBC89x
DHCPv6 Guard Overview
The DHCPv6 Guard feature blocks reply and advertisement messages that come from unauthorized DHCP servers and relay agents. Packets are classified into one of the three DHCP type messages. All client messages are always switched regardless of device role. DHCP server messages are only processed further if the device role is set to server. Further processing of server messages includes DHCP server advertisements(for source validation and server preference) and DHCP server replies (for permitted prefixes). If the device is configured as a DHCP server, all the messages need to be switched, regardless of the device role configuration.
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_dhcp/configuration/15-sy/dhcp-15-sy-book/ip6-dhcpv6-guard.pdf

upvoted 3 times

Question #431 *Topic 1*

The network administrator configured CoPP so that all SNMP traffic from Cisco Prime located at 192.168.1.11 toward the router CPU is limited to 1000 kbps. Any traffic that exceeds this limit must be dropped.

```
access-list 100 permit udp any any eq 161
!
class-map CM-SNMP
match access-group 100
!
policy-map PM-COPP
class CM-SNMP
police 1000 conform-action transmit
!
control-plane
service-policy input PM-COPP
```

The network administrator is not getting the desired result for the SNMP traffic and SNMP traffic is getting dropped frequently. Which set of configurations resolves the issue?

A. no access-list 100

```
access-list 100 permit tcp host 192.168.1.11 any eq 161
```

B. no access-list 100

```
access-list 100 permit udp host 192.168.1.11 any eq 161
!
policy-map PM-COPP
class CM-SNMP
no police 1000 conform-action transmit
police 1000000 conform-action transmit
!
control-plane
no service-policy input PM-COPP
!
interface E 0/0
service-policy input PM-COPP
!
interface E 0/1
service-policy input PM-COPP
```

C. no access-list 100

```
access-list 100 permit udp host 192.168.1.11 any eq 161
!
policy-map PM-COPP
class CM-SNMP
no police 1000 conform-action transmit
police 1000000 conform-action transmit
```

D. policy-map PM-COPP

```
class CM-SNMP
no police 1000 conform-action transmit
police 1000000 conform-action transmit
```

Correct Answer: *C*

inteldarvid 5 months, 1 week ago

Selected Answer: C

yes option C correct

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: C

Agree on "C", because it is more specific than "D". ("D" could work, too.)
1000 kbps = 1 000 000 bits per second => bps is used in CoPP, so this value needs to be changed.
The ACL syntax for matching the interesting traffic:
access-list 100 permit udp host <snmp manager> any <snmp agents> eq 161 (or eq snmp)

Example:
https://community.cisco.com/t5/routing/acl-help/td-p/2788816/page/2

upvoted 3 times

Question #432 *Topic 1*

Which protocol must be secured with MD-5 authentication across the MPLS cloud to prevent hackers from introducing bogus routers?

A. RSVP

B. ALSO

C. LDP

D. MP-BGP

Correct Answer: *C*

inteldarvid 4 months, 3 weeks ago

Selected Answer: C

yes, the give anwser is correct

upvoted 1 times

DUBC89x 1 year ago

Selected Answer: C

The Label Distribution Protocol (LDP) can also be secured with MD-5 authentication across the MPLS cloud. This prevents hackers from introducing bogus routers, which would participate in the LDP.
https://www.cisco.com/c/en/us/td/docs/net_mgmt/vpn_solutions_center/2-0/mpls/provisioning/guide/PGmpls1.html

upvoted 3 times

Question #433

Topic 1

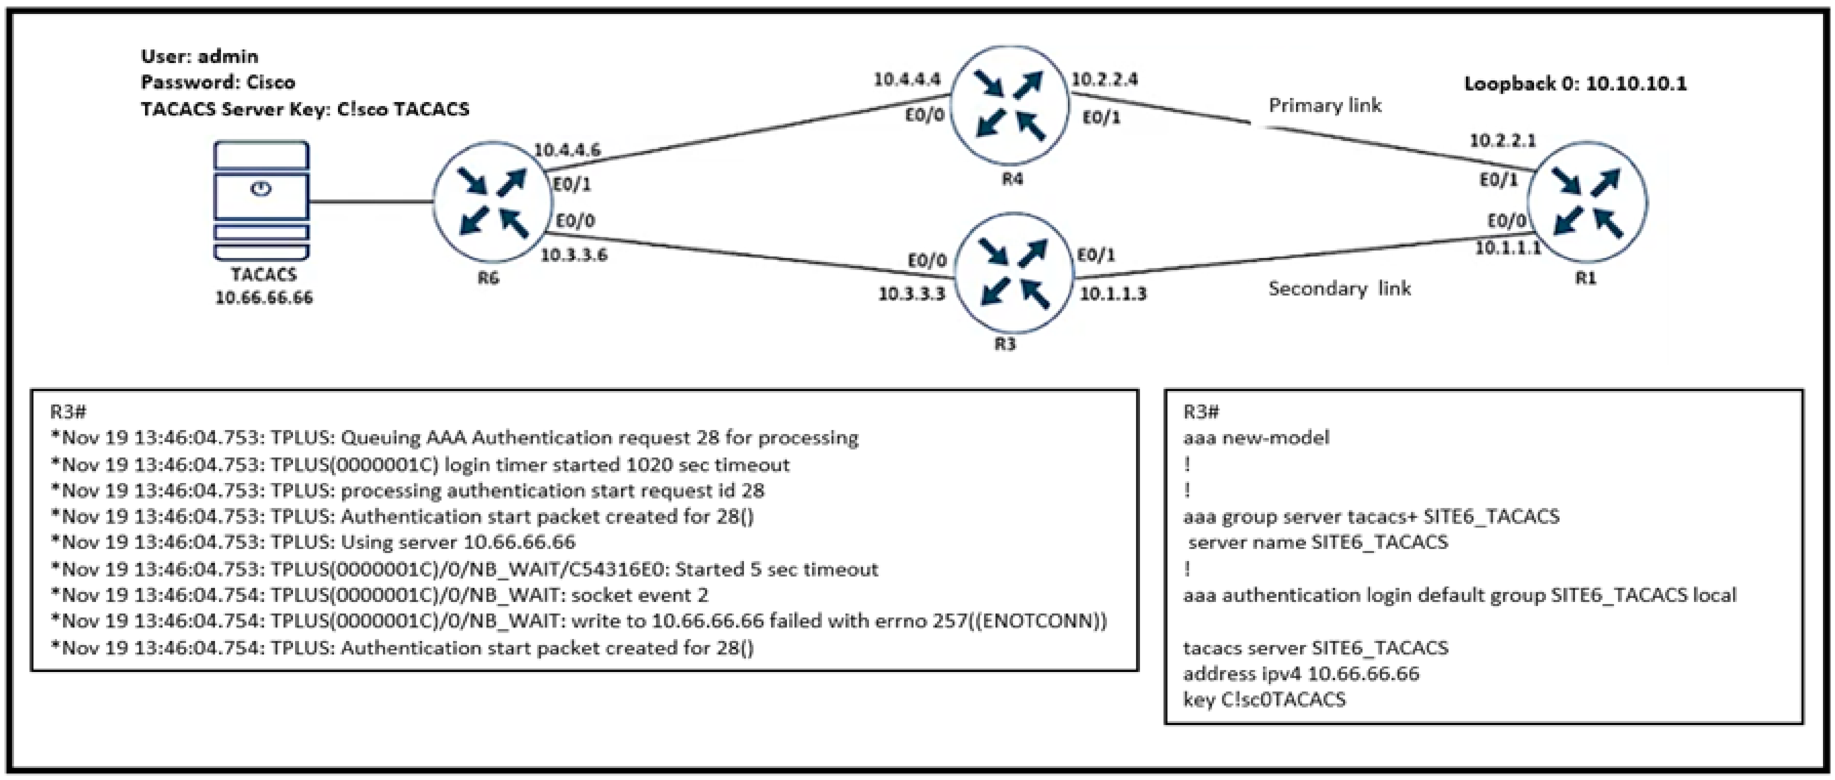

```
R3#
*Nov 19 13:46:04.753: TPLUS: Queuing AAA Authentication request 28 for processing
*Nov 19 13:46:04.753: TPLUS(0000001C) login timer started 1020 sec timeout
*Nov 19 13:46:04.753: TPLUS: processing authentication start request id 28
*Nov 19 13:46:04.753: TPLUS: Authentication start packet created for 28()
*Nov 19 13:46:04.753: TPLUS: Using server 10.66.66.66
*Nov 19 13:46:04.753: TPLUS(0000001C)/0/NB_WAIT/C54316E0: Started 5 sec timeout
*Nov 19 13:46:04.754: TPLUS(0000001C)/0/NB_WAIT: socket event 2
*Nov 19 13:46:04.754: TPLUS(0000001C)/0/NB_WAIT: write to 10.66.66.66 failed with errno 257((ENOTCONN))
*Nov 19 13:46:04.754: TPLUS: Authentication start packet created for 28()
```

```
R3#
aaa new-model
!
!
aaa group server tacacs+ SITE6_TACACS
 server name SITE6_TACACS
!
aaa authentication login default group SITE6_TACACS local

tacacs server SITE6_TACACS
address ipv4 10.66.66.66
key C!sc0TACACS
```

Refer to the exhibit R3 cannot authenticate via TACACS. Which configuration resolves the issue?

A. tacacs server SITE6_TACACS
key C!sc0TACACS

B. tacacs server SITE6_TACACS
key C!scoTACACS

C. tacacs server SITE6_TACACS
address ipv4 10.60.66.66
key C!scoTACACS

D. tacacs server SITE6_TACACS
address ipv4 10.66.66.66
key CiscoTACACS

Correct Answer: *D*

herojacky Highly Voted 11 months, 3 weeks ago

Selected Answer: B

key C!scoTACACS

upvoted 10 times

DeWalt95 Most Recent 3 weeks ago

Can someone explain why its B? Looks like the router has the correct config already?

upvoted 1 times

inteldarvid 4 months, 3 weeks ago

Selected Answer: B

yey option B is correct

upvoted 1 times

[Removed] 8 months, 1 week ago

Isn't space considered a character? if yes then none is correct.

upvoted 3 times

forccnp 9 months, 1 week ago

Selected Answer: B

B is the correct one

upvoted 2 times

Lilienen 10 months ago

Selected Answer: B

B is correct

upvoted 1 times

Dejjie 10 months, 1 week ago

Option D is wrong!

upvoted 1 times

Question #434 *Topic 1*

A customer is running an mGRE DMVPN tunnel over WAN infrastructure between hub and spoke sites. The existing configuration allows NHRP to add spoke routers automatically to the multicast NHRP mappings. The customer is migrating the network from IPv4 to the IPv6 addressing scheme for those spokes' routers that support IPv6 and can run DMVPN tunnel over the IPv6 network.

Which configuration must be applied to support IPv4 and IPv6 DMVPN tunnels on spoke routers?

A. tunnel mode ipv6ip 6to4

B. tunnel mode ipv6ip auto-tunnel

C. tunnel mode ipv6ip 6rd

D. tunnel mode ipv6ip isatap

Correct Answer: *B*

SAMAKEMM 1 month ago

Selected Answer: B

tunnel mode ipv6ip isatap does not support multicast and it is used for intra-sites

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: B

option B:
https://netcraftsmen.com/ccie-rs-prep-ipv6-part-2/

upvoted 2 times

HungarianDish 6 months, 2 weeks ago

Selected Answer: A

Quick overview:
IPv6 Automatic IPv4-Compatible Tunnels (and also point-to-multipoint):
IPv4-compatible: tunnel mode ipv6ip auto-tunnel -> not recommended either by Cisco nor by IETF/rfc7059
More details:
https://netcraftsmen.com/ccie-rs-prep-ipv6-part-2/
https://datatracker.ietf.org/doc/html/rfc7059#section-3.2
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/interface/configuration/15-s/ir-15-s-book/ip6-auto-comp-tun.pdf

6to4: tunnel mode ipv6ip 6to4 -> for site-to-site/branch-to-branch
Limitations: no IGP support, only static route or BGP

6RD: 6RD builds upon the 6to4 tunneling mechanism, and is designed for ISPs (for provider core) -> not for Branch-to-Branch

ISATAP: tunnel mode ipv6ip isatap -> designed for intrasite tunneling (within a site/branch), still can be run between sites, too
Limitations: -no IPv6 multicast, OSPFv3, EIGRP: neighbors are formed manually with "neighbor" command
or need to use static route

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

I doubt that cisco would want the answer "B" with "tunnel mode ipv6ip auto-tunnel".
I think that the best fit is "A" "tunnel mode ipv6ip 6to4", however, "D" ISATAP could be an option, too.
Question mentions the ability to "add spoke routers automatically to the multicast NHRP mappings" -> I am not sure that it would work with ISATAP.
It is hard to decide on the answer without testing dmvpn+6to4 and ISATAP. It think that such a lab is out of scope for ENARSI.

All of these are IPv6 over IPv4 tunnel technologies, but the question requests "DMVPN tunnel over the IPv6 network".
I wonder if the question is written correctly.

upvoted 1 times

forccnp 9 months ago

Selected Answer: B

tunnel mode ipv6ip auto-tunnel
Example:
Router(config-if)# tunnel mode ipv6ip
auto-tunnel

upvoted 2 times

Lilienen 10 months, 1 week ago

Selected Answer: D

Based on the document from ellen_AA, I vote for answer D

upvoted 1 times

⊟ **ellen_AA** 11 months, 1 week ago

Selected Answer: B

Given answer is correct
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/interface/configuration/15-s/ir-15-s-book/ip6-auto-comp-tun.pdf

upvoted 2 times

Question #435 *Topic 1*

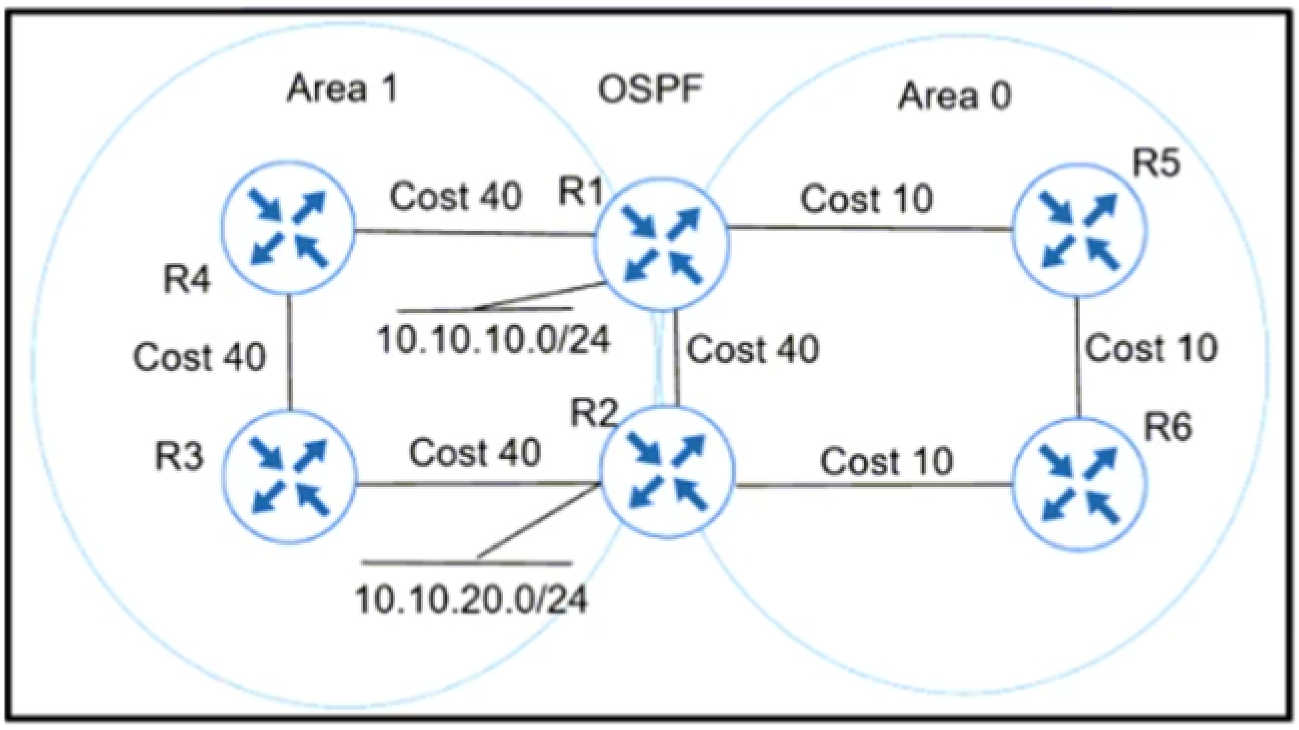

Refer to the exhibit. Which action ensures that 10.10.10.0/24 reaches 10.10.20.0/24 through the direct link between R1 and R2?

A. Configure R1 and R2 LAN links as nonpassive.

B. Configure R1 and R2 links under area 1.

C. Configure OSPF link cost to 1 between R1 and R2.

D. Configure OSPF path cost to 3 between R1 and R2.

Correct Answer: *B*

⊟ **ellen_AA** Highly Voted 11 months ago

Selected Answer: B

Both given routes are in area1, direct link is in area0. By default ospf prefers intra-area routes 'O' over inter-area routes 'O IA'. Bringing the direct link to the same area as both routes will make it preferable.

upvoted 13 times

⊟ **HungarianDish** 7 months, 1 week ago

Great solution! Well explained!
https://networklessons.com/ospf/ospf-path-selection-explained
OSPF will first look at the “type of path” to make a decision and, secondly look at the metric (cost).

upvoted 2 times

⊟ **GReddy2323** 10 months, 2 weeks ago

Thank you for your detailed answer.

upvoted 1 times

⊟ **DeWalt95** Most Recent 3 weeks ago

Selected Answer: B

OSPF route prerence.
Intra-Area, Inter-Area, external.
Takes precedence over metric

upvoted 2 times

⊟ **smayus** 10 months ago

How do we know if they are not in the same area 1 for sure? Wouldn't that be easier to configure cost to 3 for example?

upvoted 1 times

⊟ **abd123** 10 months, 2 weeks ago

why not C , please advise

upvoted 1 times

Question #436 Topic 1

The summary route is not shown in the Router_B routing table after this below configuration on Router_A:

```
interface ethernet 0
description location ID:S4318T3E77F02
ip address 192.168.3.1 255.255.255.0
ip summary-address eigrp 1 172.16.80.0 255.255.240.0
```

Which Router_A configuration resolves the issue by advertising the summary route to Router_B?

A.
```
interface loopback 0
ip address 172.18.81.1 255.255.255.0
interface Ethernet 0
ip address 192.168.3.1 255.255.255.0
ip summary-address eigrp 1 172.16.80.0 255.255.240.0
```

B.
```
interface loopback 0
ip address 172.16.79.1 255.255.255.0
interface Ethernet 0
ip address 192.168.3.1 255.255.255.0
ip summary-address eigrp 1 172.16.80.0 255.255.240.0
```

C.
```
interface loopback 0
ip address 172.16.81.1 255.255.255.0
interface Ethernet 0
ip address 192.168.3.1 255.255.255.0
ip summary-address eigrp 1 172.16.80.0 255.255.240.0
```

D.
```
interface loopback 0
ip address 172.16.96.1 255.255.255.0
interface Ethernet 0
ip address 192.168.3.1 255.255.255.0
ip summary-address eigrp 1 172.16.80.0 255.255.240.0
```

Correct Answer: *C*

DeWalt95 3 weeks ago

Selected Answer: C

Will only advertise a summary if there is valid address within it in the RIB.

C is the only address in the summary range.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

correct

upvoted 1 times

Titini 10 months, 1 week ago

We need to add a loopback interface that belongs to the subnet 172.16.80.0 255.255.240.0

upvoted 3 times

GReddy2323 10 months, 2 weeks ago

I don't understand this question, can someone please explain?

upvoted 1 times

targetA 10 months ago

Summary addresses won't show in the routing table unless there's a componant route in the summary address range

Another way of asking this question would be "which of the following loopback ip addresses are contained in the 172.16.80.0 255.255.240.0 summary address"

A - 172.18.81.1 - is not in the summary as it has to be 172.16.80.0 - 172.16.95.255

B - 172.16.79.1 - is just under the summary
D - 172.16.96.1 - is just over the summary

Only C is in that range, which triggers the summary address to go into the routing table
upvoted 9 times

Question #437 *Topic 1*

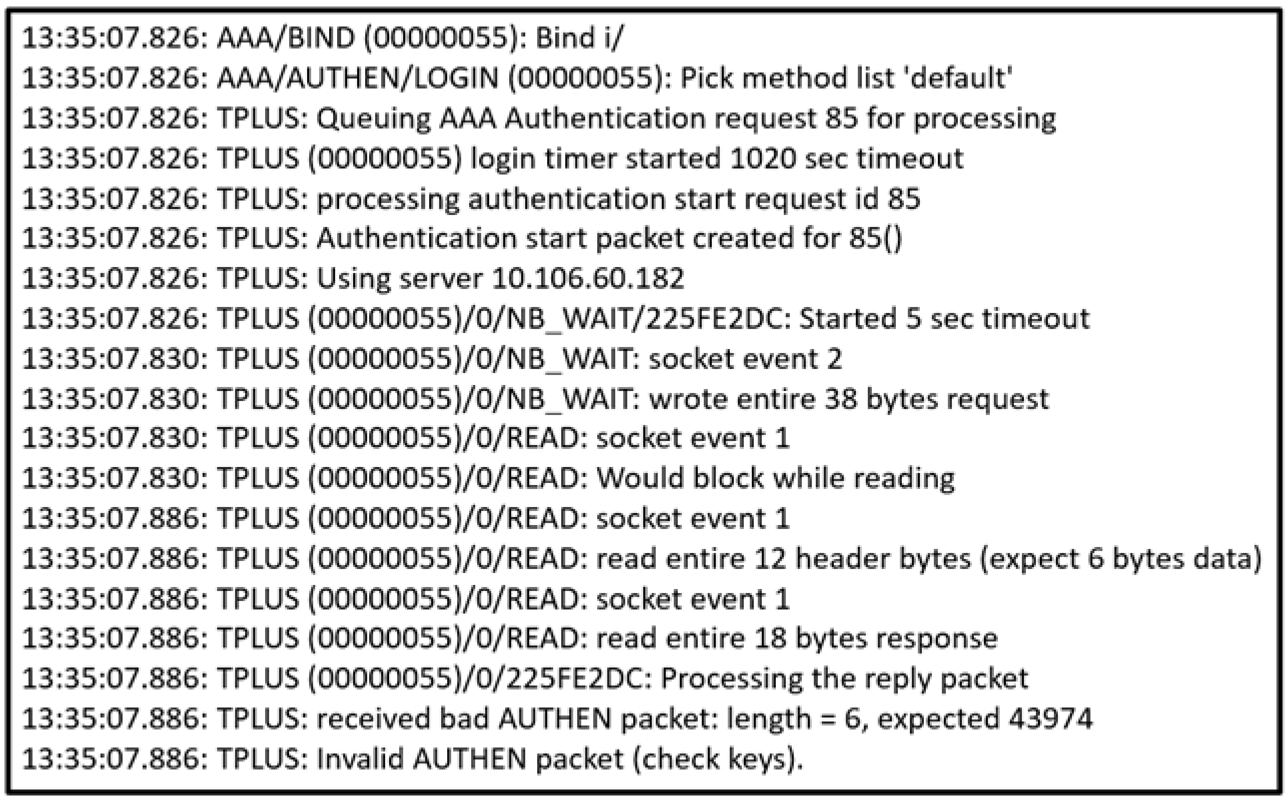

```
13:35:07.826: AAA/BIND (00000055): Bind i/
13:35:07.826: AAA/AUTHEN/LOGIN (00000055): Pick method list 'default'
13:35:07.826: TPLUS: Queuing AAA Authentication request 85 for processing
13:35:07.826: TPLUS (00000055) login timer started 1020 sec timeout
13:35:07.826: TPLUS: processing authentication start request id 85
13:35:07.826: TPLUS: Authentication start packet created for 85()
13:35:07.826: TPLUS: Using server 10.106.60.182
13:35:07.826: TPLUS (00000055)/0/NB_WAIT/225FE2DC: Started 5 sec timeout
13:35:07.830: TPLUS (00000055)/0/NB_WAIT: socket event 2
13:35:07.830: TPLUS (00000055)/0/NB_WAIT: wrote entire 38 bytes request
13:35:07.830: TPLUS (00000055)/0/READ: socket event 1
13:35:07.830: TPLUS (00000055)/0/READ: Would block while reading
13:35:07.886: TPLUS (00000055)/0/READ: socket event 1
13:35:07.886: TPLUS (00000055)/0/READ: read entire 12 header bytes (expect 6 bytes data)
13:35:07.886: TPLUS (00000055)/0/READ: socket event 1
13:35:07.886: TPLUS (00000055)/0/READ: read entire 18 bytes response
13:35:07.886: TPLUS (00000055)/0/225FE2DC: Processing the reply packet
13:35:07.886: TPLUS: received bad AUTHEN packet: length = 6, expected 43974
13:35:07.886: TPLUS: Invalid AUTHEN packet (check keys).
```

Refer to the exhibit. Which action resolves the authentication problem?

A. Configure the same password between the TACACS+ server and router.

B. Configure the TCP port 49 to be reachable by the router.

C. Configure the UDP port 1812 to be allowed on the TACACS+ server.

D. Configure the user name on the TACACS+ server.

Correct Answer: *A*

⊟ **jansan55** 2 months, 3 weeks ago

Selected Answer: A

Reference: https://www.cisco.com/c/en/us/support/docs/security-vpn/terminal-access-controller-access-control-system-tacacs-/200467-Troubleshoot-TACACS-Authentication-Issue.html
This is a sample debug output from the Router, when the TACACS server is configured with a wrong pre shared key.
upvoted 1 times

⊟ **HarwinderSekhon** 4 months ago

last line of debug "check keys" we must match keys/passwords
upvoted 1 times

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: A

correct A
upvoted 1 times

Question #438 *Topic 1*

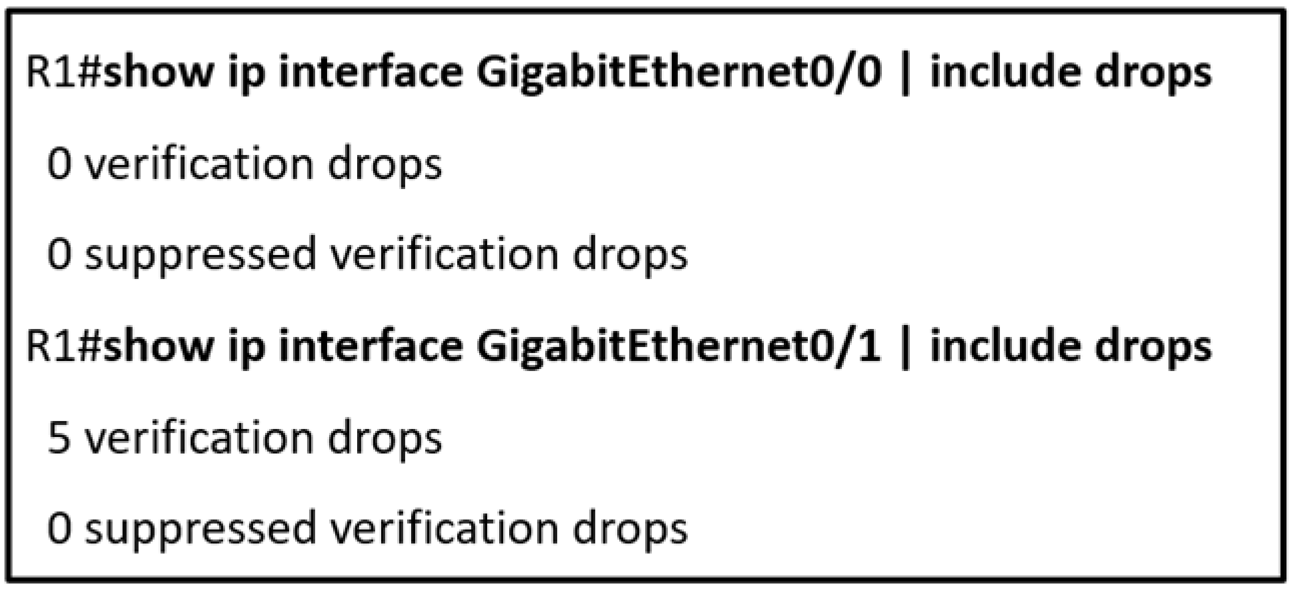

```
R1#show ip interface GigabitEthernet0/0 | include drops
 0 verification drops
 0 suppressed verification drops
R1#show ip interface GigabitEthernet0/1 | include drops
 5 verification drops
 0 suppressed verification drops
```

Refer to the exhibit. R1 is configured with uRPF, and ping to R1 is failing from a source present in the R1 routing table via the GigabitEthernet 0/0 interface. Which action resolves the issue?

A. Enable Cisco Express Forwarding to ensure that uRPF is functioning correctly.

B. Modify the uRPF mode from strict to loose.

C. Add a floating static route to the source on R1 to the GigabitEthernet0/1 interface.

D. Remove the access list from the interface GigabitEthernet 0/0.

Correct Answer: *B*

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: B

yes correct is B

upvoted 2 times

⊟ **Titini** 10 months, 1 week ago

Selected Answer: B

with uRPF strict mode, this can cause legitimate packets to be dropped, so that's something to keep in mind. In uRPF loose mode, the source of the IP packet must simply appear in the routing table. So loose mode looks for any interface in the routing table other than the default route.

upvoted 4 times

⊟ **ellen_AA** 11 months, 1 week ago

I think it's A, if it's set to strict mode we would have seen drops at the interface. Because uRPF strict mode verifies reachability of the source + its existing in the routing table.

upvoted 1 times

Question #439 Topic 1

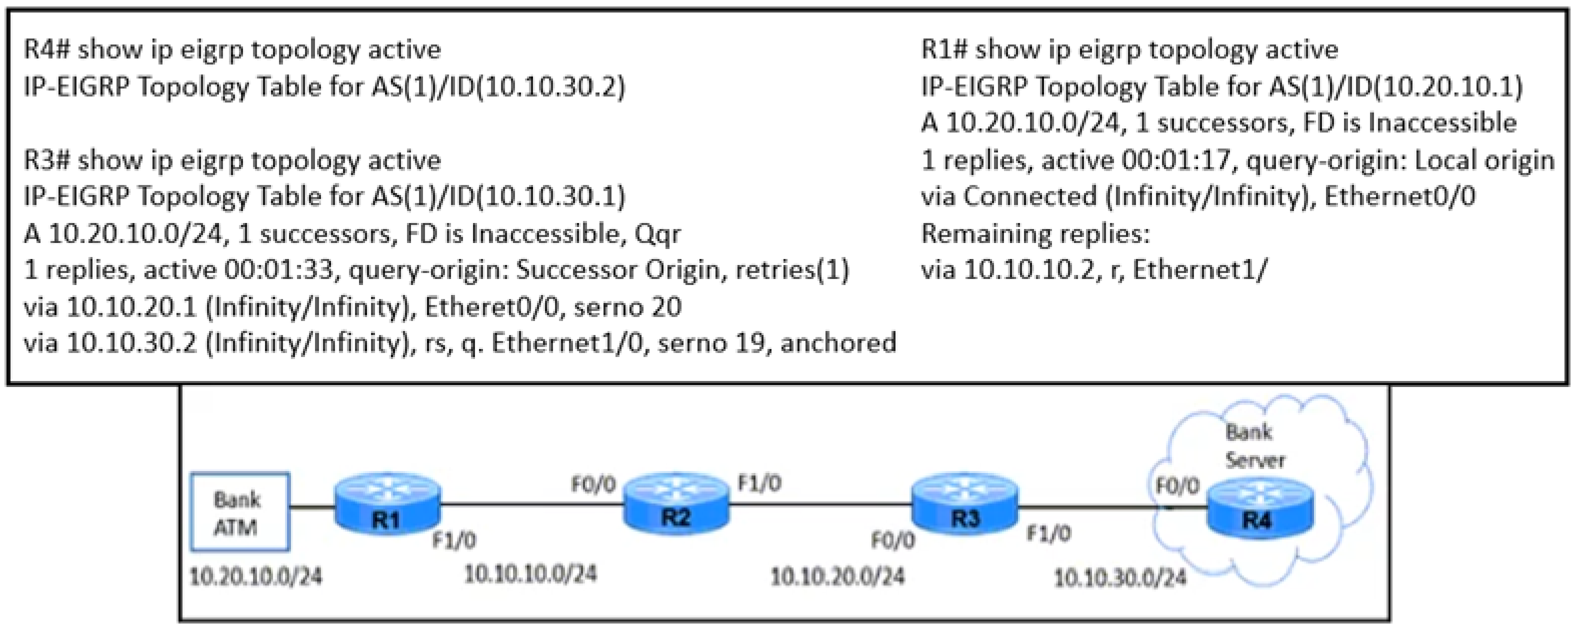

```
R4# show ip eigrp topology active
IP-EIGRP Topology Table for AS(1)/ID(10.10.30.2)

R3# show ip eigrp topology active
IP-EIGRP Topology Table for AS(1)/ID(10.10.30.1)
A 10.20.10.0/24, 1 successors, FD is Inaccessible, Qqr
1 replies, active 00:01:33, query-origin: Successor Origin, retries(1)
via 10.10.20.1 (Infinity/Infinity), Etheret0/0, serno 20
via 10.10.30.2 (Infinity/Infinity), rs, q. Ethernet1/0, serno 19, anchored
```

```
R1# show ip eigrp topology active
IP-EIGRP Topology Table for AS(1)/ID(10.20.10.1)
A 10.20.10.0/24, 1 successors, FD is Inaccessible
1 replies, active 00:01:17, query-origin: Local origin
via Connected (Infinity/Infinity), Ethernet0/0
Remaining replies:
via 10.10.10.2, r, Ethernet1/
```

Refer to the exhibit. A bank ATM site has difficulty connecting with the bank server. A network engineer troubleshoots the issue and finds that R4 has no active route to the bank ATM site. Which action resolves the issue?

A. EIGRP peering between R1 and R2 to be fixed.

B. Advertise 10.10.30.0/24 subnet in R3 EIGRP AS.

C. Advertise 10.10.30.0/24 subnet in R1 EIGRP AS.

D. EIGRP peering between R3 and R4 to be fixed.

Correct Answer: *B*

louisvuitton12 1 month, 2 weeks ago

Selected Answer: D

Check previous docs

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: D

option "D" correct:
https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/BRKRST-2331.pdf

upvoted 1 times

SolidSnake74 6 months ago

Answer is D
B & C are false, just because no reason to work on that route because R1 is already missing the Bank ATM network
If you advertise a new network (10.10.30.0/24), it won't change anything about the connectivity issue
From R1 output, R1 lost the 10.20.10.0/24 and is searching for a feasible successor
At least the Query reached R3 as it says the Query is comming from the successor origin
In output of R3, 'via 10.10.30.2', "rs" means retry & SIA => this means didn't get a repy within the 3 minutes timer of the SIA (stuck-in-active)
Last point, the fact that R4 gives no output, just means not execting to search anymore the a route or still not notified by the query from R3
So, R3 is pending on retry and is in SIA state, because no answer from R4, then R4 has no route in active state because it has no being queried and still doesn't known a route is missing or it already removed it

upvoted 2 times

HungarianDish 6 months, 3 weeks ago

The LAN on R1 goes unreachable. => R1 starts querying for a path for prefix 10.20.10.0/24. (R1 is active for this prefix.)
After this, the routers R2 and R3 also go active and query.
"show ip eigrp topology active" -> Look for the small r in the command output.
This means that the router awaits a reply to a query for that prefix from that neighbor.
The R3 also goes active and awaits a reply from R4.
As shown, the router R4 does not go active for the prefix, because it did not recieve any queries. So, there must be an issue with the link between R3 and R4.
For same reason, R3 does not receive an SIA reply from R4, so the "s" appears for that neighbor under "show ip eigrp topology active".
rs = r - reply Status, s - sia Status
sent query and waiting for reply + did not recieve reply in SIA time

upvoted 3 times

HungarianDish 6 months, 3 weeks ago

https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/BRKRST-2331.pdf
There are no active routes on other end (R4)

https://community.cisco.com/t5/routing/eigrp-stuck-in-active-routes/td-p/478346
if EIGRP has had a route in its table (some prefix that it can route to)
and if EIGRP loses that prefix from the table then EIGRP should put that prefix into "active state".

If a neighbor does not respond within a certain limit of time the EIGRP router assumes that it may have lost sync with the router and tears down the neighbor relationship and starts over again from the beginning.

upvoted 3 times

⊟ **HungarianDish** 6 months, 3 weeks ago

R1 waiting for reply about prefix 10.20.10.0/24 from neighbor 10.10.10.2
IP-EIGRP Topology Table for AS(1)/ID(10.20.10.1)
via 10.10.10.2, r, Ethernet1/

"Look for the small r in the command output. This means that the router awaits a reply to a query for that prefix from that neighbor."
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.html

you must discover the location where one or more routers sends queries and does not receive replies, while the downstream router is not in this state.
For example, the router could send queries and they are acknowledged, but the reply from the downstream router is not received.

the router R4 did not go active for the prefix, so the problem must be between routers R3 and R4.
After some time, we see that R3 kills the neighborship to R4 and declares an SIA state:

The router R3 also sends SIA queries to R4, but it does not receive an SIA reply from R4.

Once the router sends an SIA query but does not receive an SIA reply, the s appears for that neighbor

upvoted 2 times

⊟ **HungarianDish** 7 months, 1 week ago

Selected Answer: D

For me it is answer "D". Please, see:
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.html
https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/BRKRST-2331.pdf

upvoted 2 times

⊟ **Typovy** 8 months ago

None is correct. Having an route with active status is not good in eigrp. None will resolve issue since R1 has this route already marked as active. R1 need to resolve connected network problem.

upvoted 2 times

⊟ **Lilienen** 10 months ago

D. EIGRP peering between R3 and R4 to be fixed.
R3 has routes, R4 has no routes, therefore there's no peering between R3 and R4

upvoted 4 times

⊟ **ellen_AA** 11 months, 1 week ago

None is correct, I think we should advertise 10.20.10.0/24 into R1 correctly, so it wouldn't show as unreachable. Closer answer is C but with the correct subnet

upvoted 2 times

Question #440 *Topic 1*

Refer to the exhibit.

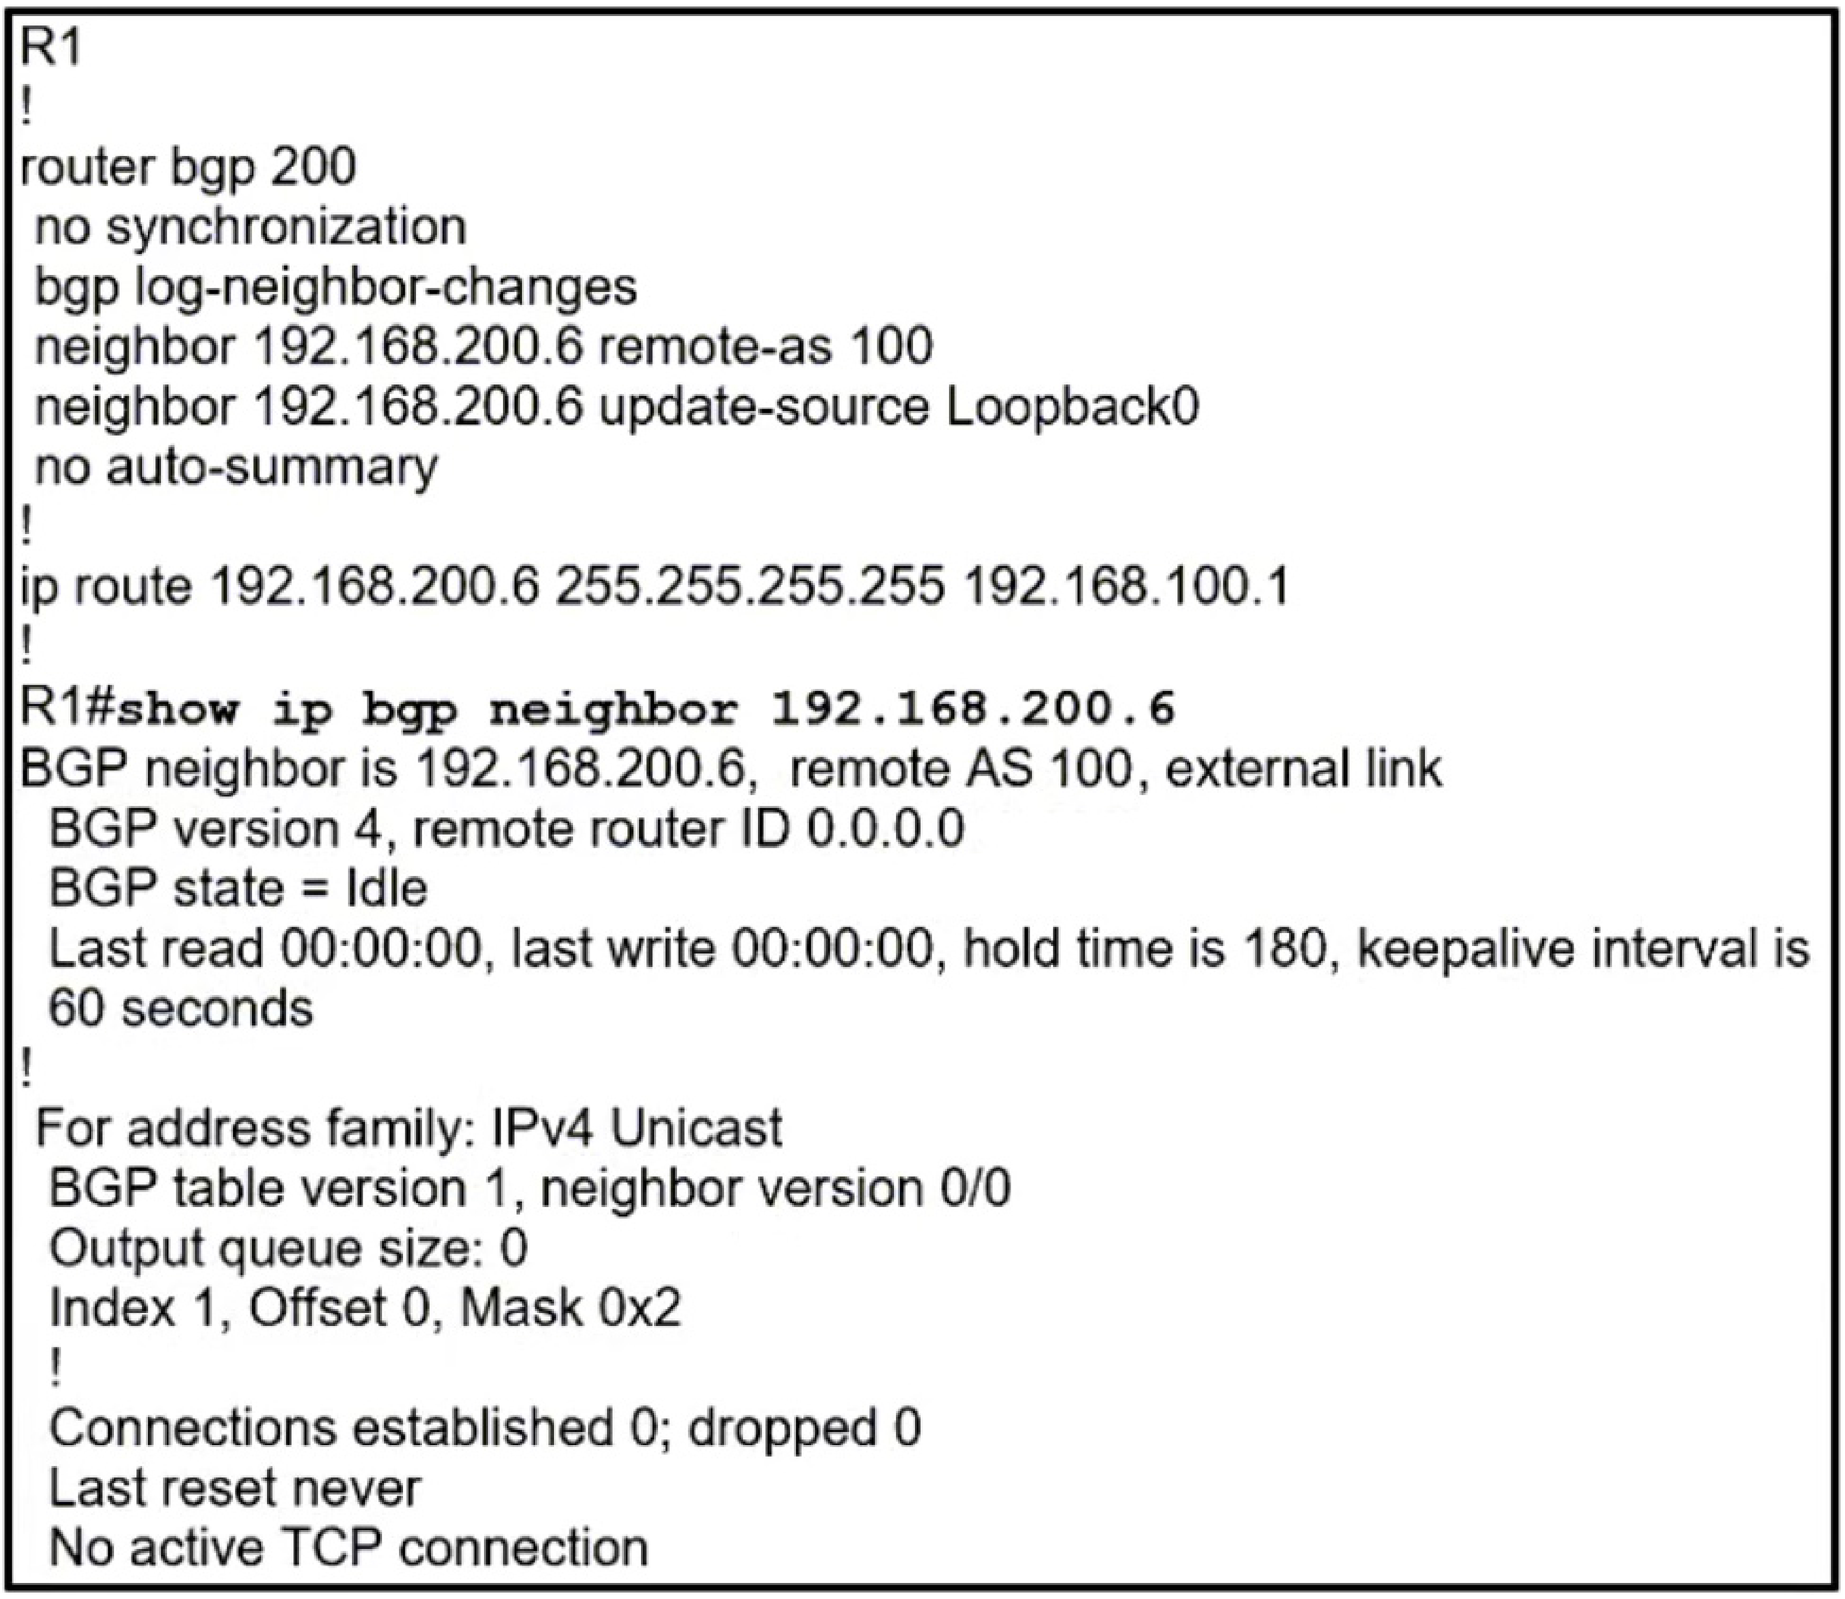

```
R1
!
router bgp 200
 no synchronization
 bgp log-neighbor-changes
 neighbor 192.168.200.6 remote-as 100
 neighbor 192.168.200.6 update-source Loopback0
 no auto-summary
!
ip route 192.168.200.6 255.255.255.255 192.168.100.1
!
R1#show ip bgp neighbor 192.168.200.6
BGP neighbor is 192.168.200.6,  remote AS 100, external link
  BGP version 4, remote router ID 0.0.0.0
  BGP state = Idle
  Last read 00:00:00, last write 00:00:00, hold time is 180, keepalive interval is
  60 seconds
!
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 0/0
  Output queue size: 0
  Index 1, Offset 0, Mask 0x2
  !
  Connections established 0; dropped 0
  Last reset never
  No active TCP connection
```

The BGP neighbor is not coming up. Which action resolves the issue?

A. Configure the ebgp-multihop 2 command on R1 toward the neighbor.

B. Configure a valid router ID on the neighbor that shows an invalid router ID of 0.0.0.0.

C. The route map on eBGP sessions must allow the prefixes from the neighbor.

D. Enable synchronization between the neighbors to bring the neighborship up.

Correct Answer: *A*

Question #441 *Topic 1*

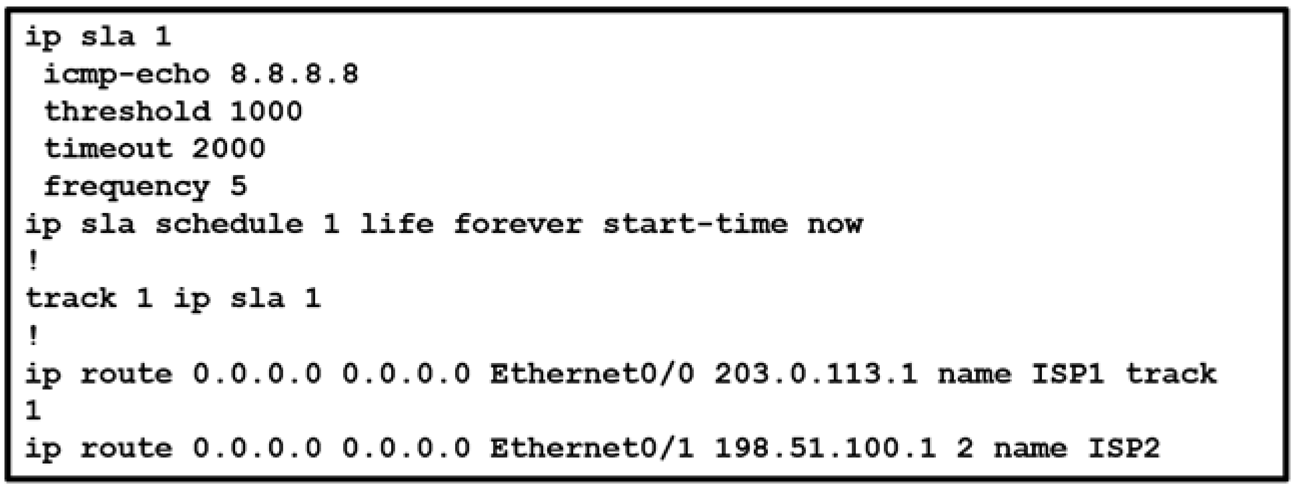

```
ip sla 1
 icmp-echo 8.8.8.8
 threshold 1000
 timeout 2000
 frequency 5
ip sla schedule 1 life forever start-time now
!
track 1 ip sla 1
!
ip route 0.0.0.0 0.0.0.0 Ethernet0/0 203.0.113.1 name ISP1 track
1
ip route 0.0.0.0 0.0.0.0 Ethernet0/1 198.51.100.1 2 name ISP2
```

Refer to the exhibit. After recovering from a power failure, Ethernet0/1 stayed down while Ethernet0/0 returned to the up/up state. The default route through ISP1 was not reinstated in the routing table until Ethernet0/1 also came up.

Which action resolves the issue?

A. Add a static route to the 8.8.8.8/32 destination through the next hop 203.0.113.1.

B. Remove the references to the interface names from both static default routes.

C. Reference the track object 1 in both static default routes.

D. Configure the default route through ISP1 with a higher administrative distance than 2.

Correct Answer: *C*

Zizu007 Highly Voted 11 months, 3 weeks ago

Selected Answer: A

IP SLA operation to check ISP-1 should not rely on ISP-2 interface/connection--reach 8.8.8.8 via ISP-2 in this case.
in this case either;
- the address for icmp-echo should be changed to 203.0.113.1
OR
- Add static route 8.8.8.8/32 -->203.0.113.1

upvoted 6 times

[Removed] Most Recent 4 months, 3 weeks ago

I initially did not understand this scenario, and I had to create a lab to see the behavior.

Here is what is happening.
Router reboots, and link E0/1 remains down. While the router reboots, it is attempting to rebuild the RIB, but because there is a track object monitoring the 0/0 via E0/0 to 203.0.113.1, that route is never installed into the RIB due to the SLA failing probes to 8.8.8.8/32. By adding a separate static route to 8.8.8.8 via E0/0, the SLA successfully probes 8.8.8.8 and the default route is installed into the RIB.

upvoted 2 times

inteldarvid 5 months, 1 week ago

Selected Answer: A

option A correct

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

Selected Answer: A

Confirmed "A" in lab. Of course, with the original configuration in the output, the track object (+sla) did not go down at all, when lost connectivity to 203.0.113.1, because the target of the sla (8.8.8.8) was still available through 198.51.100.1. Solution "A" is definitely needed for the whole thing to work. It would be best to add this line, too.
ip route 8.8.8.8 255.255.255.255 Null0 2
https://community.cisco.com/t5/routing/ip-sla-tracking-a-far-ip/td-p/1971337

upvoted 2 times

Typovy 9 months, 1 week ago

Selected Answer: A

A is correct

upvoted 2 times

Question #442

Topic 1

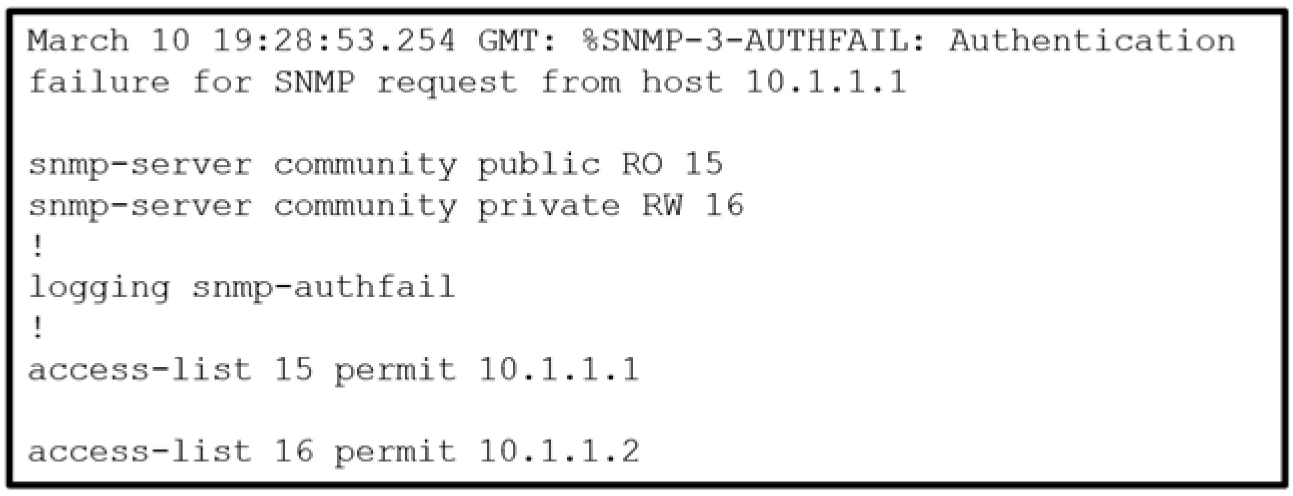

```
March 10 19:28:53.254 GMT: %SNMP-3-AUTHFAIL: Authentication
failure for SNMP request from host 10.1.1.1

snmp-server community public RO 15
snmp-server community private RW 16
!
logging snmp-authfail
!
access-list 15 permit 10.1.1.1

access-list 16 permit 10.1.1.2
```

Refer to the exhibit. Which action resolves the issue?

A. Configure host IP address in access-list 16.

B. Configure SNMPv3 on the router.

C. Configure SNMP authentication on the router.

D. Configure a valid SNMP community string.

Correct Answer: *D*

Titini Highly Voted 10 months ago

Selected Answer: D

Looking at the configuration, it seems that there are two SNMP community strings configured, "public" and "private". However, the configuration does not specify which community string is being used by the host with IP address 10.1.1.1.

The correct action to resolve this issue would be to configure a valid SNMP community string on the device from which the SNMP request is being made, and ensure that it matches the community string configured on the router.

upvoted 8 times

Brand Most Recent 3 months, 2 weeks ago

"Which action resolves the issue?" where exactly? On the router? Because we can do it by adding the 10.1.1.1 to ACL16. But if it's going to be on the 10.1.1.1 we can do it by matching the community.

upvoted 1 times

guy276465281819372 4 months, 3 weeks ago

Selected Answer: A

A seems more correct

upvoted 2 times

Almylle 5 months, 3 weeks ago

Selected Answer: D

Refer this page pls https://www.routerfreak.com/what-to-do-about-snmp-3-authfail-messages/

upvoted 1 times

Brand 3 months, 2 weeks ago

snmp-server community public RO 5
snmp-server community private RW 6
!
! log incorrect SNMP Communities
!
logging snmp-authfail
!
! access-list for RO
!
access-list 5 permit 10.1.1.100
access-list 5 permit 10.1.1.101
!
! access-list for RW
!
access-list 6 permit 10.1.1.101

This is the configuration in the link and according to this we need to add the 10.1.1.1 to the ACL 16 to give the guy access for RW as well.

I'd go for A

upvoted 1 times

upp3r 8 months ago

Selected Answer: D

https://www.routerfreak.com/what-to-do-about-snmp-3-authfail-messages/

this explains an incorrect community string or blocking ACL both result in the same IOS log message

as we can see the source ip address in the log message it must mean the used community string is incorrect

upvoted 4 times

forccnp 9 months, 1 week ago

Selected Answer: A

I vote for A

upvoted 2 times

Lilienen 10 months ago

Selected Answer: A

I vote for A

upvoted 1 times

Question #443 *Topic 1*

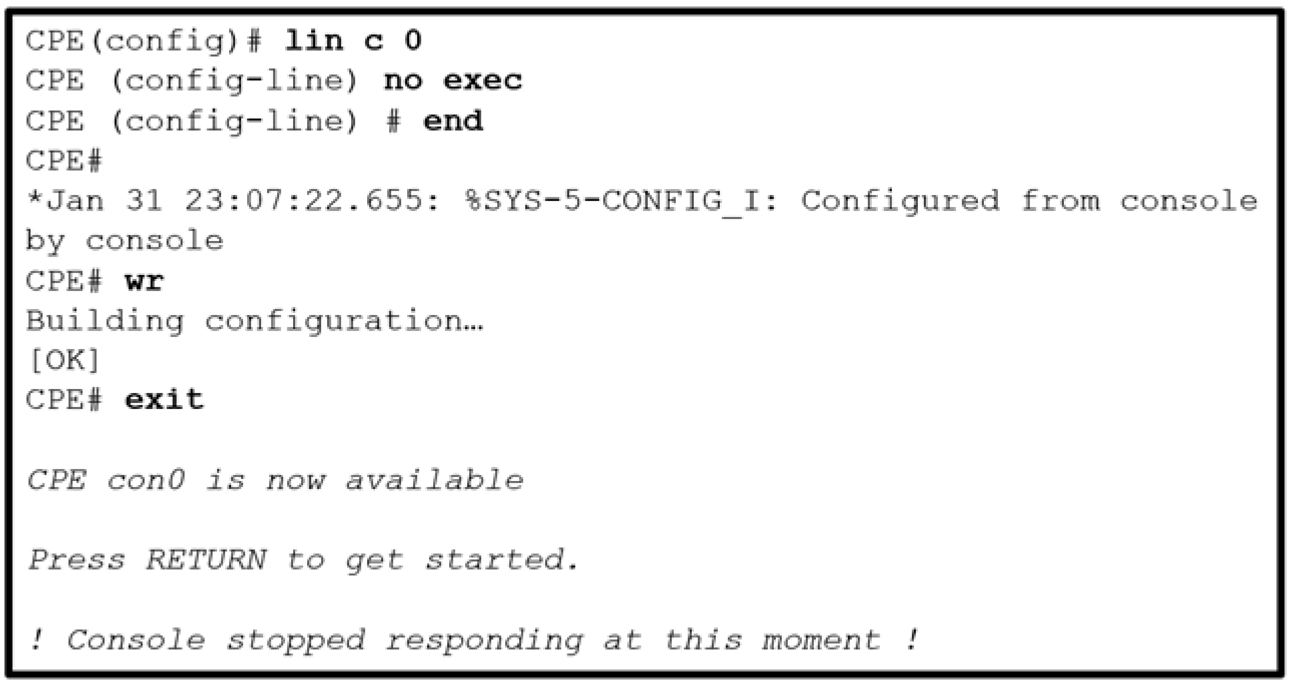

```
CPE(config)# lin c 0
CPE (config-line) no exec
CPE (config-line) # end
CPE#
*Jan 31 23:07:22.655: %SYS-5-CONFIG_I: Configured from console
by console
CPE# wr
Building configuration…
[OK]
CPE# exit

CPE con0 is now available

Press RETURN to get started.

! Console stopped responding at this moment !
```

Refer to the exhibit. An administrator is attempting to disable the automatic logout after a period of inactivity. After logging out, the console stopped responding to all keyboard inputs. Remote access through SSH still works.

Which action resolves the issue?

A. Configure the no exec-timeout command on line con 0.

B. Configure the absolute-timeout command on line con 0.

C. Configure the exec command on line con 0.

D. Configure the default exec-timeout command on line con 0

Correct Answer: *A*

dq28 Highly Voted 11 months, 3 weeks ago

Selected Answer: C

Why should configuring the exec-timeout fix the disabled exec? I would vote C to "resolv the issue" of console stopped responding to all keyboards.

https://community.cisco.com/t5/routing/no-exec/td-p/3715737

upvoted 6 times

mitosenoriko 11 months, 2 weeks ago

That's not what I'm talking about.

upvoted 2 times

HungarianDish Most Recent 7 months, 1 week ago

Selected Answer: C

As the issue is that they disabled Cisco IOS exec under the console line, I go with "C".
https://rednectar.net/2011/08/27/never-use-the-no-exec-timeout-command/
(The console session time-out is not a real issue.)

upvoted 2 times

Malasxd 7 months, 1 week ago

Selected Answer: C

What I need to fix here? The timeout or the disabled console port?

"A" if you need to fix the timeout
"C" if you need to reenable the console port.

I would go for C

upvoted 2 times

Titini 10 months ago

Selected Answer: A

We seem to have 2 issues. The first is to disable automatic logout and the second seem to be the re enablement of the exec for console line. Since we still have access via vty (hopefully and exec via vty) we can login via vty and configure option A.

upvoted 2 times

ellen_AA 11 months, 1 week ago

It's recommended not to use: #no exec-timeout
A couldn't be an answer! #exec-timeout 0 (is recommended to disable logout this way)

upvoted 1 times

Question #444 *Topic 1*

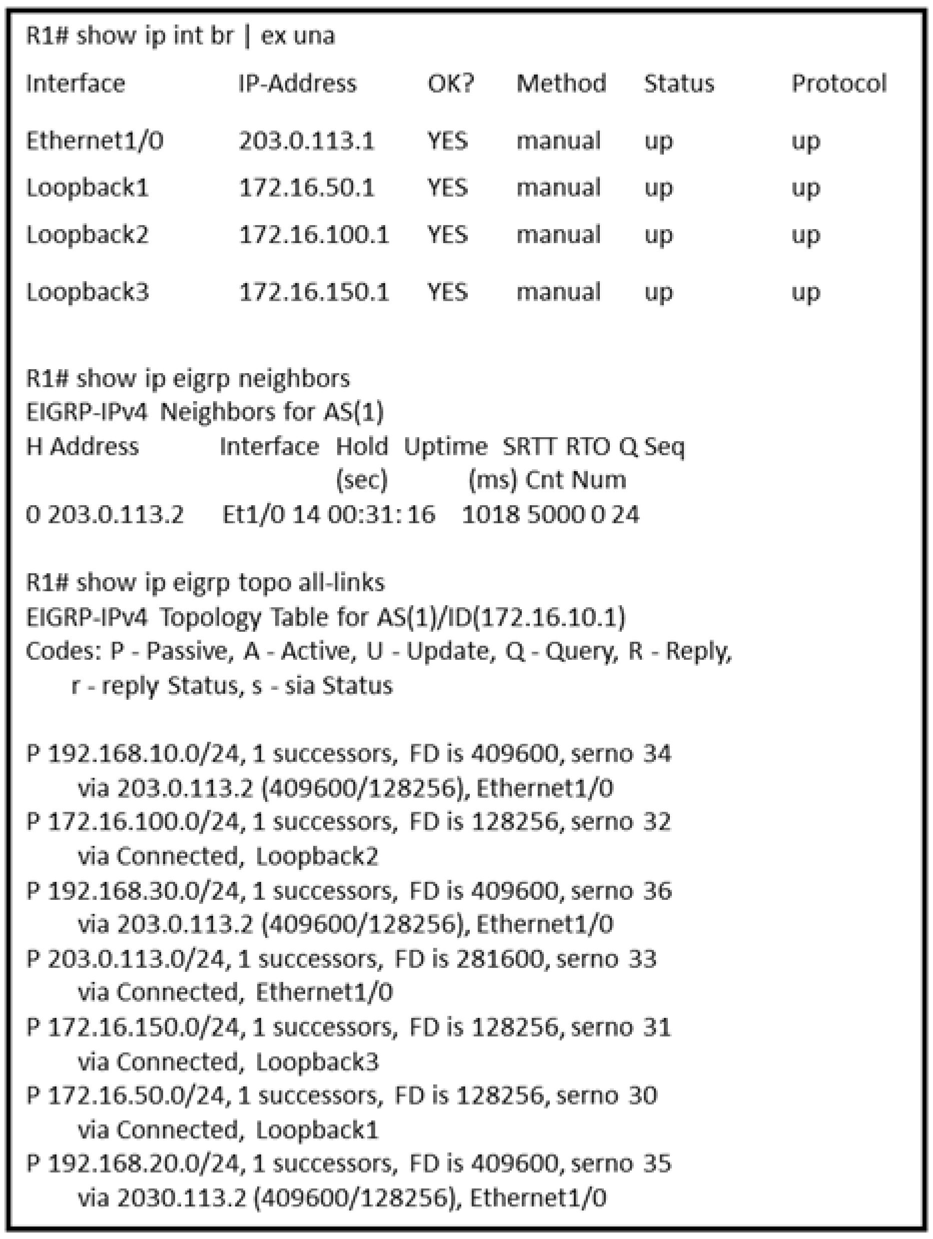

```
R1# show ip int br | ex una
Interface        IP-Address     OK?   Method    Status        Protocol
Ethernet1/0      203.0.113.1    YES   manual    up            up
Loopback1        172.16.50.1    YES   manual    up            up
Loopback2        172.16.100.1   YES   manual    up            up
Loopback3        172.16.150.1   YES   manual    up            up

R1# show ip eigrp neighbors
EIGRP-IPv4  Neighbors for AS(1)
H Address          Interface  Hold  Uptime  SRTT RTO Q Seq
                              (sec)         (ms) Cnt Num
0 203.0.113.2      Et1/0 14 00:31:16   1018 5000 0 24

R1# show ip eigrp topo all-links
EIGRP-IPv4  Topology Table for AS(1)/ID(172.16.10.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 192.168.10.0/24, 1 successors,  FD is 409600, serno 34
        via 203.0.113.2 (409600/128256), Ethernet1/0
P 172.16.100.0/24, 1 successors,  FD is 128256, serno 32
        via Connected,  Loopback2
P 192.168.30.0/24, 1 successors,  FD is 409600, serno 36
        via 203.0.113.2 (409600/128256), Ethernet1/0
P 203.0.113.0/24, 1 successors,  FD is 281600, serno 33
        via Connected,  Ethernet1/0
P 172.16.150.0/24, 1 successors,  FD is 128256, serno 31
        via Connected,  Loopback3
P 172.16.50.0/24, 1 successors,  FD is 128256, serno 30
        via Connected,  Loopback1
P 192.168.20.0/24, 1 successors,  FD is 409600, serno 35
        via 2030.113.2 (409600/128256), Ethernet1/0
```

Refer to the exhibit. Routers R1 and R2 have established a network adjacency using EIGRP, and both routers are advertising subnets to its neighbor. After issuing the show ip EIGRP topology all-links command in R1, some prefixes are not showing R2 as a successor. Which action resolves the issue?

A. Configure the network statement on the neighbor.

B. Rectify the incorrect router ID in R2.

C. Resolve the incorrect metric on the link.

D. Enable split-horizon.

Correct Answer: *C*

SolidSnake74 Highly Voted 6 months ago

Answer is A
Can't be B, because if it has been the case, no route at all from R2 would have been learned
Can't be C, because a metric mismatch leads in no neighbor relationship and we have
Can't be D, Split horizon is enabled by default and it just avoid sharing the network via a same interface it received. It means i can GET routes from R2 but not just sharing them back to R2 via the same interface.

upvoted 8 times

Youssefmetry 4 months, 1 week ago

EIGRP require the K values to match only .. not bandwidth or delay. I think C is a possible answer.

upvoted 1 times

fizzer Most Recent 3 months, 1 week ago

Selected Answer: B

Correct answer is B

I have just labbed this up, two EIGRP routers with the same router ID, they did not accept each others locally generated routes, however all other EIGRP learned routes that did not originate from them were accepted

It does make sense for the receiving router to feel weird receiving a route with its own router ID as the originating router, as soon as I changed the router ID on one of them, they both accepted each others locally generated routes (installed them in the topology table), changed router ID to match again and the routes disappeared again.

The reason some routes shows up according to the question is because those routes were not generated by R2

upvoted 1 times

[Removed] 4 months, 3 weeks ago

Selected Answer: A

as explained by solidsnake74. The best answer is A.

B, the router ID is different, otherwise no routes would be learned at all, neighborship forms, but no routes learned.
C, no neighborship with an error message pointing out K-value mismatch
D, no.

upvoted 1 times

daloslav 6 months, 3 weeks ago

Selected Answer: B

If two routers have the same router ID, they can become neighbors but some prefixes (specifically external redistributed prefixes) will be missing.

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

router id R2: 203.0.113.2 (from "sh ip eigrp neigh")
router id R1: 172.16.10.1 (from "sh ip eigrp top all")
The router ids are different.

upvoted 1 times

yellowswan 3 months, 2 weeks ago

203.0.113.2 is interface address, not RID

upvoted 1 times

Gedson 5 months ago

router id R2: 203.0.113.2, Doesn't rourter ID it's interface's ip

upvoted 1 times

Malasxd 7 months, 2 weeks ago

I didn't get it. The routes do not has R2 as sucessors are directly connected. You cannot change it.

upvoted 2 times

HungarianDish 7 months, 1 week ago

Agree. I do not see any incorrect metrics either. I tried to figure it out in the lab, but information about the topology is missing, and thus I can't conclude to any of the answers.
In some topology, not all routes are listed as feasible successors due to split-horizont. (Still we would not disable SH just because of that.)
https://community.cisco.com/t5/switching/eigrp-topology-table-all-links/td-p/3179536

upvoted 2 times

Xerath 9 months ago

Selected Answer: C

I think "C" is the best answer.

upvoted 2 times

Lilienen 10 months, 1 week ago

Selected Answer: A

I vote for A

upvoted 1 times

abd123 10 months, 2 weeks ago

why? please

upvoted 2 times

Question #445 *Topic 1*

An engineer configures PBR on R5 and wants to create a policy that matches traffic destined toward 10.10.10.0/24 and forwards it toward 10.1.1.1. This traffic must also have its IP precedence set to 5. All other traffic should be forwarded toward 10.1.1.2 and have its IP precedence set to 0. Which configuration meets the requirements?

```
A. access-list 100 permit ip any 10.10.10.0 0.0.0.255
route-map CCNP permit 10
match ip address 100
set ip next-hop 10.1.1.1
set ip precedence 5
!
route-map CCNP permit 20
set ip next-hop 10.1.1.2
set ip precedence 0
```

```
B. access-list 100 permit ip any 10.10.10.0 0.0.0.255
route-map CCNP permit 10
match ip address 100
set ip next-hop 10.1.1.1
set ip precedence 0
!
route-map CCNP permit 20
set ip next-hop 10.1.1.2
set ip precedence 5
!
route-map CCNP permit 30
```

```
C. access-list 1 permit 10.10.10.0 0.0.0.255
route-map CCNP permit 10
match ip address 1
set ip next-hop 10.1.1.1
set ip precedence 5
!
route-map CCNP permit 20
set ip next-hop 10.1.1.2
set ip precedence 0
```

```
D. access-list 1 permit 10.10.10.0 0.0.0.255
access-list 2 permit any
route-map CCNP permit 10
match ip address 1
set ip next-hop 10.1.1.1
set ip precedence 5
!
route-map CCNP permit 20
match ip address 2
set ip next-hop 10.1.1.2
set ip precedence 0
!
route-map CCNP permit 30
```

Correct Answer: *A*

[Removed] 4 months, 3 weeks ago

Selected Answer: A

Correct, you need to define a source and destination for the Access-list, and two route maps to match each of the parameters for that destination and to be manipulated for the next-hop

upvoted 2 times

HungarianDish 7 months, 1 week ago

Selected Answer: A

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2/25ew/configuration/guide/conf/pbroute.pdf
Classify traffic based on extended access list criteria.

upvoted 2 times

Question #446

Topic 1

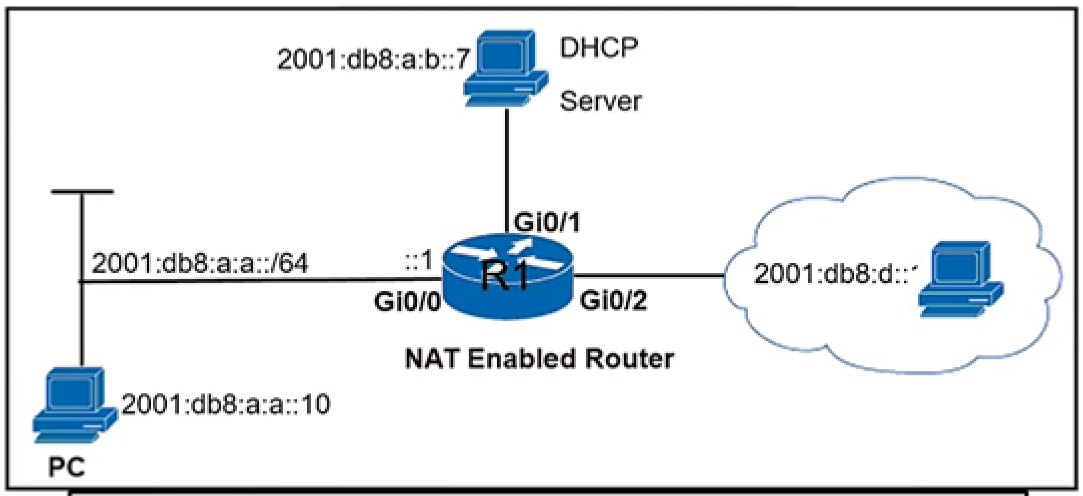

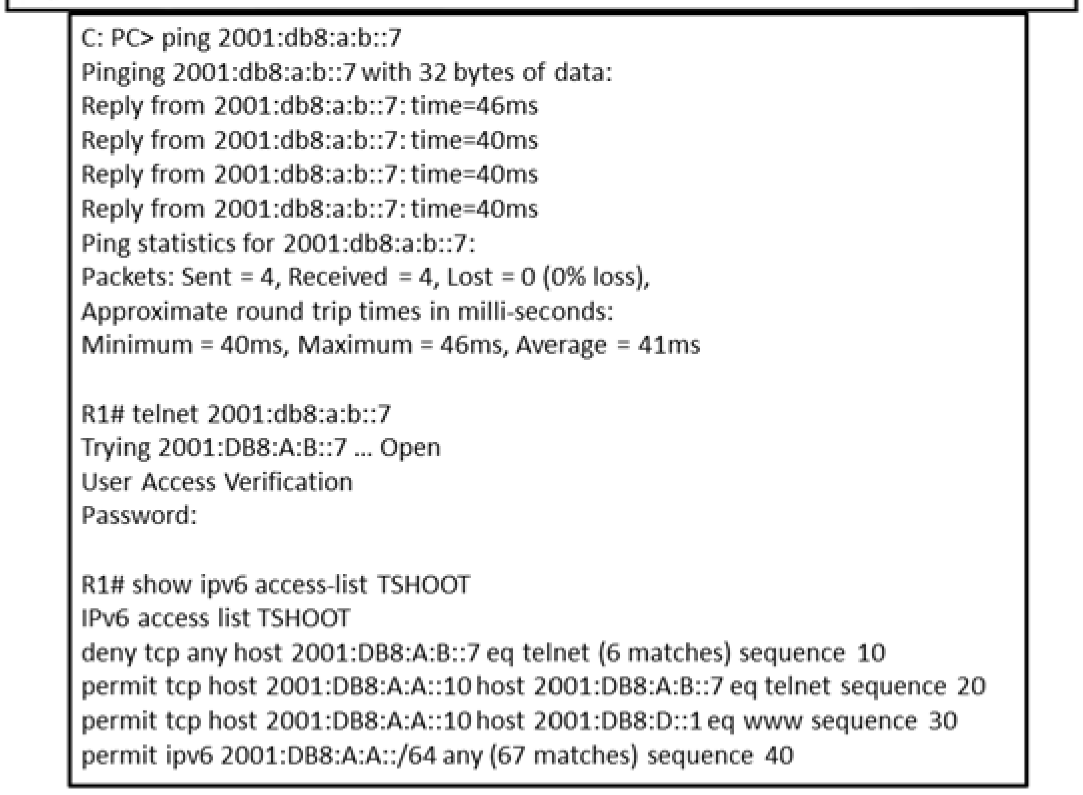

```
C: PC> ping 2001:db8:a:b::7
Pinging 2001:db8:a:b::7 with 32 bytes of data:
Reply from 2001:db8:a:b::7: time=46ms
Reply from 2001:db8:a:b::7: time=40ms
Reply from 2001:db8:a:b::7: time=40ms
Reply from 2001:db8:a:b::7: time=40ms
Ping statistics for 2001:db8:a:b::7:
Packets: Sent = 4, Received  = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 40ms, Maximum = 46ms, Average = 41ms

R1# telnet 2001:db8:a:b::7
Trying 2001:DB8:A:B::7 ... Open
User Access Verification
Password:

R1# show ipv6 access-list TSHOOT
IPv6 access list TSHOOT
deny tcp any host 2001:DB8:A:B::7 eq telnet (6 matches) sequence 10
permit tcp host 2001:DB8:A:A::10 host 2001:DB8:A:B::7 eq telnet sequence 20
permit tcp host 2001:DB8:A:A::10 host 2001:DB8:D::1 eq www sequence 30
permit ipv6 2001:DB8:A:A::/64 any (67 matches) sequence 40
```

Refer to the exhibit. An engineer is troubleshooting a failed Telnet session from PC to the DHCP server. Which action resolves the issue?

A. Remove sequence 30 and add it back to the IPv6 traffic filter as sequence 5.

B. Remove sequence 20 for sequence 40 in the access list to allow Telnet.

C. Remove sequence 10 to add the PC source IP address and add it back as sequence 10.

D. Remove sequence 20 and add it back to the IPv6 traffic filter as sequence 5.

Correct Answer: *D*

inteldarvid 4 months, 3 weeks ago

Selected Answer: D

the give answer is correct

upvoted 1 times

Question #447 Topic 1

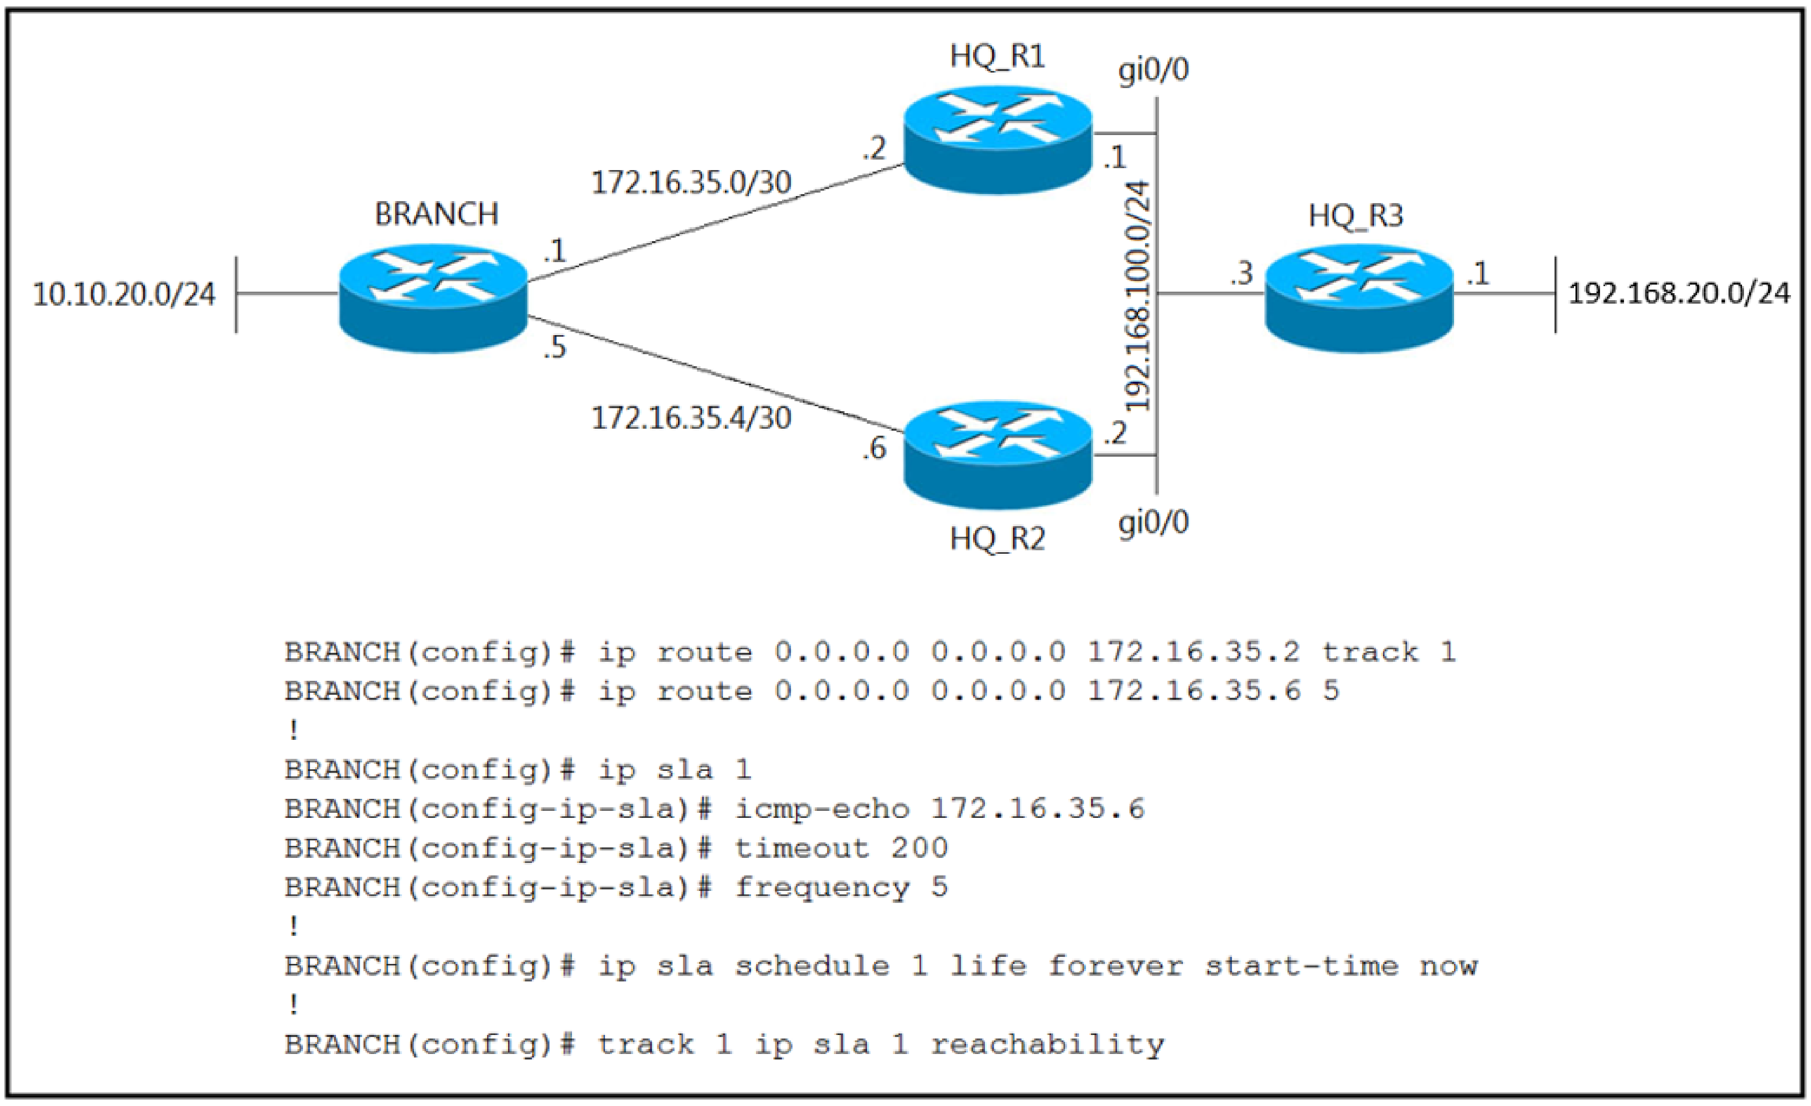

```
BRANCH(config)# ip route 0.0.0.0 0.0.0.0 172.16.35.2 track 1
BRANCH(config)# ip route 0.0.0.0 0.0.0.0 172.16.35.6 5
!
BRANCH(config)# ip sla 1
BRANCH(config-ip-sla)# icmp-echo 172.16.35.6
BRANCH(config-ip-sla)# timeout 200
BRANCH(config-ip-sla)# frequency 5
!
BRANCH(config)# ip sla schedule 1 life forever start-time now
!
BRANCH(config)# track 1 ip sla 1 reachability
```

Refer to the exhibit. Traffic from the branch network should route through HQ_R1 unless the path is unavailable. An engineer tests this functionality by shutting down interface on the BRANCH router toward HQ_R1 router, but 192.168.20.0/24 is no longer reachable from the branch router. Which set of configurations resolves the issue?

A. HQ_R1(config)# ip sla responder -
HQ_R1(config)# ip sla responder icmp-echo 172.16.35.2

B. BRANCH(config)# ip sla 1 -
BRANCH(config-ip-sla)# icmp-echo 172.16.35.1

C. HQ_R2(config)# ip sla responder -
HQ_R2(config)# ip sla responder icmp-echo 172.16.35.5

D. BRANCH(config)# ip sla 1 -
BRANCH(config-ip-sla)# icmp-echo 172.16.35.2

Correct Answer: *D*

[Removed] 4 months, 3 weeks ago

Selected Answer: D

Took me a minute to see what is going on.

The IP SLA is pointing to HQ_R2 address, and using that address (172.16.35.6) to track the static route pointing to HQ_R1. When the engineer shutdown the interface between Branch and HQ_R1, the IP SLA is still functioning and the static route to HQ_R1 remains up, black-holing traffic in the process. Change the IP SLA to the HQ_R1

upvoted 1 times

inteldarvid 4 months, 3 weeks ago

Selected Answer: D

THE GIVE ANSWER IS CORRECT

upvoted 1 times

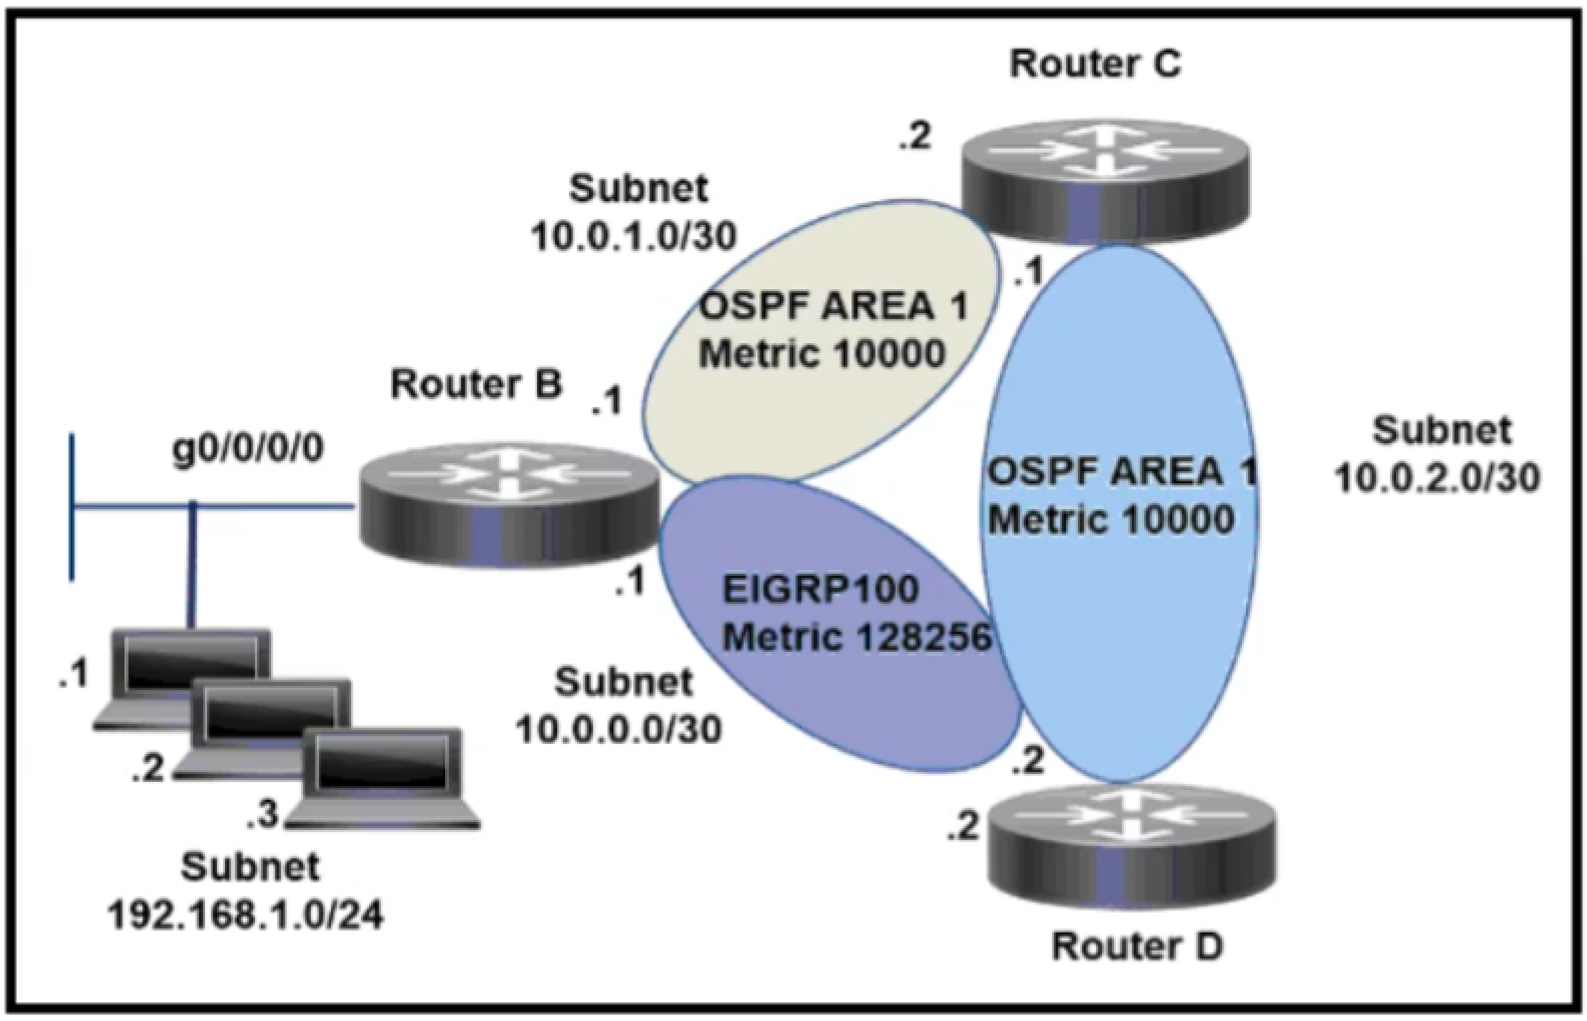

Refer to the exhibit. An engineer configures router B to direct all traffic from host 192.168.1.3 to router C. All other traffic must be routed through normal routing-protocol operations. Which configuration accomplishes the task?

A. interface g0/0/0
ip address 192.168.1.254 255.255.255.0
!
access-list 101 permit ip host 192.168.1.3 any
access-list 101 permit ip any any
!
route-map CCNP permit 10
match ip address 101
set ip next-hop 10.0.1.2

B. interface g0/0/0
ip address 192.168.1.254 255.255.255.0
ip policy route-map CCNP
!
access-list 101 permit ip host 192.168.1.3 any
!
route-map CCNP permit 10
match ip address 101
set ip next-hop 10.0.2.1

C. interface g0/0/0
ip address 192.168.1.254 255.255.255.0
ip policy route-map CCNP
!
access-list 101 permit ip any host 192.168.1.3
!
route-map CCNP permit 10
match ip address 101
set ip next-hop 10.0.1.2

D. interface g0/0/0
ip address 192.168.1.254 255.255.255.0
ip policy route-map CCNP

```
!
access-list 101 permit ip host 192.168.1.3 any
!
route-map CCNP permit 10
match ip address 101
set ip next-hop 10.0.1.2
```

Correct Answer: *D*

inteldarvid 4 months, 3 weeks ago

Selected Answer: D

the give answer is correct

upvoted 1 times

jansan55 4 months, 3 weeks ago

A: missing the ip policy route-map CCNP under int Gi0/0
B: wrong next-hop (10.0.2.1 instead of 10.0.1.2)
C: the direction is opposite in access-list 101
D: correct

upvoted 2 times

jansan55 4 months, 2 weeks ago

Well, next-hop in answer B also point to router C. Not clear the route redistribution of 192.168.1.0/24 prefix (to EIGRP or to OSPF ASBR) So answer B possible solution too.

upvoted 1 times

GReddy2323 7 months ago

Can someone elaborate? It's either A or D and I thought it was A because of the permit ip any any statement on the ACL.

upvoted 1 times

cir_ 6 months, 3 weeks ago

The any any statement in the ACL will make all traffic route via 10.0.1.2 rather than just traffic sourced from 192.168.1.3

upvoted 3 times

Question #449 *Topic 1*

Which protocol supports labeled paths between PE routers in an MPLS network?

A. LDP

B. RSVP

C. MP-BGP

D. IGP

Correct Answer: *A*

Fenix7 3 months, 1 week ago

Yes, C is the answer!

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: C

Multi-Protocol BGP extensions carry VPN policies from PE to PE.
https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/BRKMPL-1100.pdf

upvoted 3 times

HungarianDish 7 months, 1 week ago

You can use BGP to distribute routes and MPLS labels.
https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/e_fscscl.html

upvoted 4 times

Wooker 8 months, 1 week ago

Selected Answer: C

The answer is C.

In an MPLS network, Service Provider Edge (PE) routers use Multiprotocol Border Gateway Protocol (MP-BGP) to exchange label and route information with each other. MP-BGP is the routing protocol used to create and distribute labels between service provider edge routers on an MPLS network.

upvoted 4 times

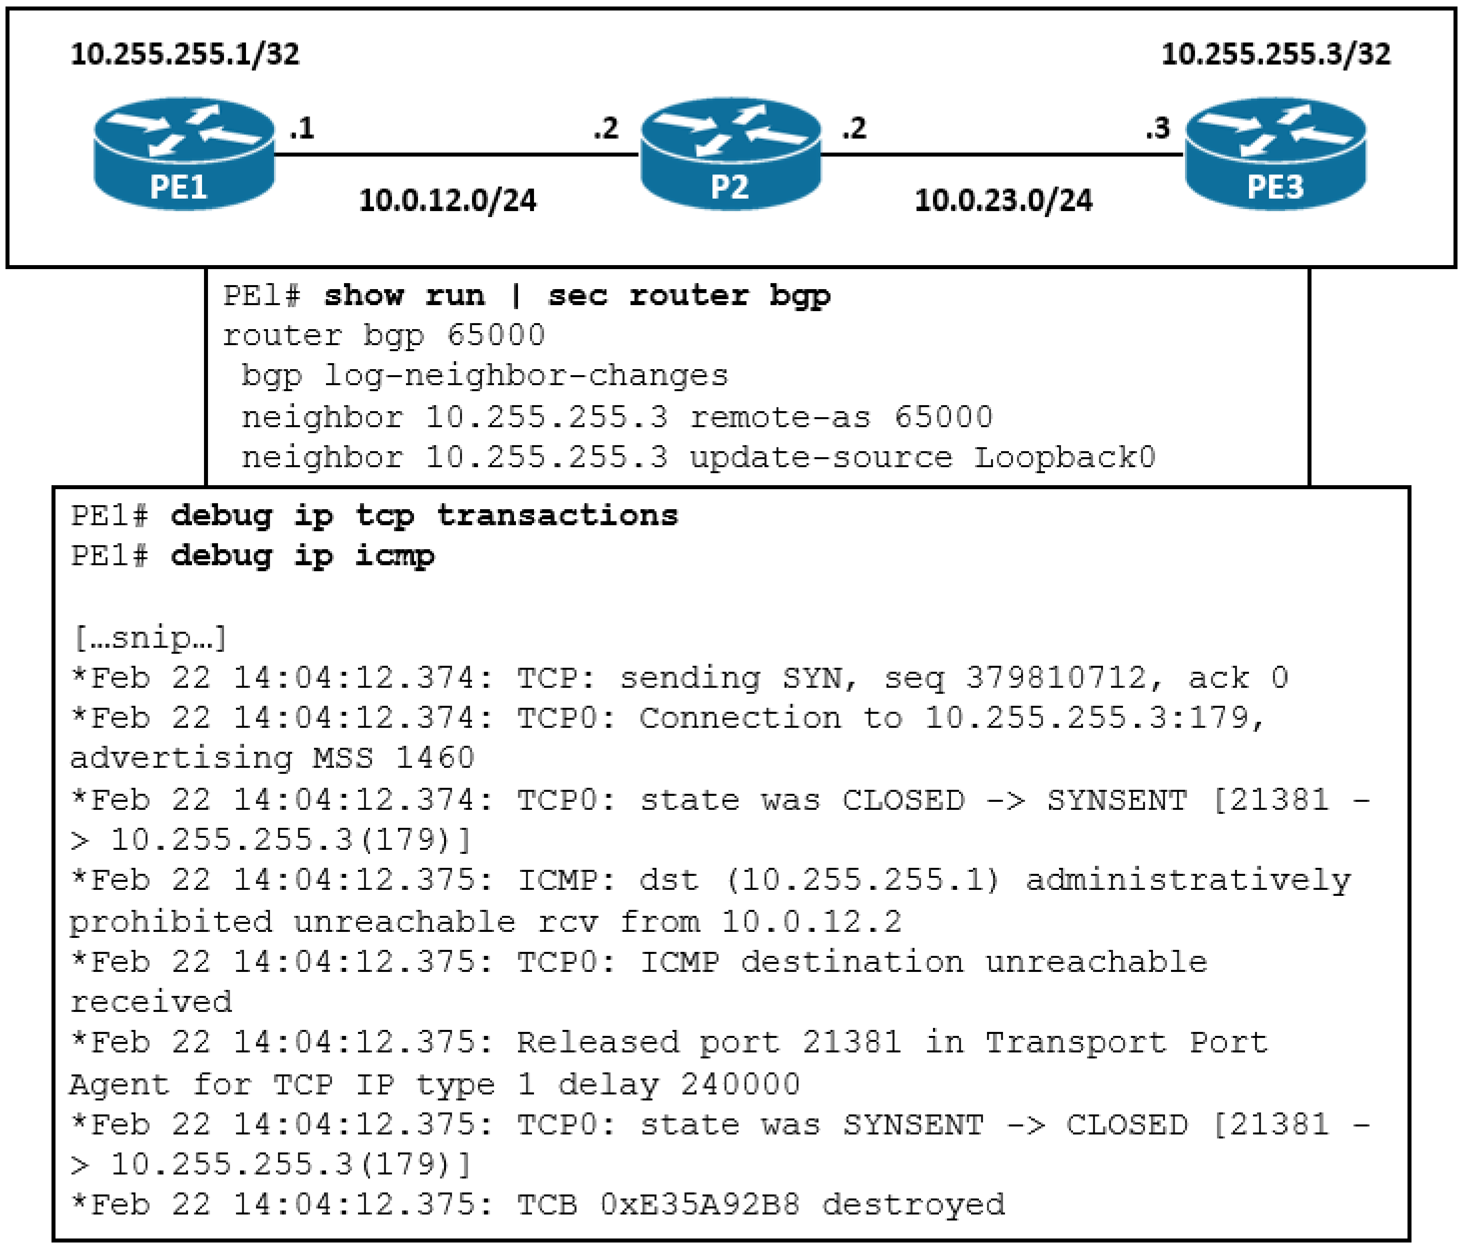

```
PE1# show run | sec router bgp
router bgp 65000
 bgp log-neighbor-changes
 neighbor 10.255.255.3 remote-as 65000
 neighbor 10.255.255.3 update-source Loopback0
```

```
PE1# debug ip tcp transactions
PE1# debug ip icmp

[…snip…]
*Feb 22 14:04:12.374: TCP: sending SYN, seq 379810712, ack 0
*Feb 22 14:04:12.374: TCP0: Connection to 10.255.255.3:179,
advertising MSS 1460
*Feb 22 14:04:12.374: TCP0: state was CLOSED -> SYNSENT [21381 -
> 10.255.255.3(179)]
*Feb 22 14:04:12.375: ICMP: dst (10.255.255.1) administratively
prohibited unreachable rcv from 10.0.12.2
*Feb 22 14:04:12.375: TCP0: ICMP destination unreachable
received
*Feb 22 14:04:12.375: Released port 21381 in Transport Port
Agent for TCP IP type 1 delay 240000
*Feb 22 14:04:12.375: TCP0: state was SYNSENT -> CLOSED [21381 -
> 10.255.255.3(179)]
*Feb 22 14:04:12.375: TCB 0xE35A92B8 destroyed
```

Refer to the exhibit. The administrator is troubleshooting a BGP peering between PE1 and PE3 that is unable to establish. Which action resolves the issue?

A. Disable sending ICMP unreachables on P2 to allow PE1 to establish a session with PE3.

B. P2 must have a route to PE3 to establish a BGP session to PE1.

C. Remove the traffic filtering rules on P2 blocking the BGP communication between PE1 and PE3.

D. Ensure that the PE3 loopback address is used as a source for BGP peering to PE1.

Correct Answer: *C*

HungarianDish Highly Voted 7 months, 1 week ago

Selected Answer: D

PE1 is trying to use PE3 loopback address for peering, so "D" is really important in this case.
"C" is unrelated to BGP. "debug ip icmp" shows administratively prohibited message for ICMP from R2. Only for ICMP and not for TCP. ICMP is unrelated to the BGP TCP process.
One more thing, they are not directly connected, so may need to enable multihop.

upvoted 8 times

ZamanR Most Recent 1 week, 4 days ago

D is the Answer

upvoted 1 times

DeWalt95 3 weeks ago

Selected Answer: D

I believe its D.

Debug ICMP is turned on and is confusing the messages..but the TCP error messages just show the TCP session timing out indicating a routing issue.

upvoted 1 times

Ghauri777 2 months ago

Selected Answer: C

Should be C. neighborship still comes up without update-source loopback command on PE3. "Administratively prohibited unreachable" message is generated when acl is applied.

upvoted 1 times

yefrimart 2 months, 1 week ago

Selected Answer: C

I labed it. I placed an ACL on PE2 blocking tcp port 179, and the logs obtains were the same, including the "ICMP destination unreachable" log, even if the ACL is not blocking the ICMP protocol itself.

upvoted 2 times

chaocheng 4 months, 1 week ago

Ans:C
lab test

```
P2#sh access-list
Extended IP access list 100
10 deny tcp host 10.255.255.1 host 10.255.255.3 eq bgp log
11 deny tcp any any eq bgp log
20 permit ip any any
```

upvoted 2 times

[Removed] 4 months, 3 weeks ago

Selected Answer: C

I'll go with C, the important thing to note in the logs is that it is "Administratively prohibited" meaning that an ACL is somehow blocking the TCP session from reaching P3 from P1

upvoted 3 times

ttl2000 5 months ago

it has to be C. if D, icmp prohibited message should from 10.0.3.0

upvoted 1 times

ttl2000 5 months ago

typo. 10.0.23.x , A R3 IP address.

upvoted 1 times

SolidSnake74 5 months ago

Answer is C
Tested in LAB each line is exactly the same and it was logical.
In the question, the "rcv from is the P2 interface", not any of the PE3 ip

```
*Jul 25 19:26:42.589: TCP: sending SYN, seq 956756274, ack 0
*Jul 25 19:26:42.589: TCP0: Connection to 1.1.1.1:179, advertising MSS 1460
*Jul 25 19:26:42.589: TCP0: state was CLOSED -> SYNSENT [54184 -> 1.1.1.1(179)]
*Jul 25 19:26:42.590: ICMP: dst (8.8.8.8) administratively prohibited unreachable rcv from 50.50.50.2
*Jul 25 19:26:42.590: TCP0: ICMP destination unreachable received
*Jul 25 19:26:42.590: Released port 54184 in Transport Port Agent for TCP IP type 1 delay 240000
*Jul 25 19:26:42.590: TCP0: state was SYNSENT -> CLOSED [54184 -> 1.1.1.1(179)]
*Jul 25 19:26:42.590: TCB 0xF6773FC0 destroyed
```

I added an ACL inbound on P2 (link between PE1 and P2) denying bgp port 179

upvoted 2 times

inteldarvid 5 months, 1 week ago

Selected Answer: D

option "D "is correct. I test in my lab:

```
R3(config-router)#do show i
*Jul 20 13:19:00.095: Reserved port 0 in Transport Port Agent for TCP IP type 0
*Jul 20 13:19:00.099: TCP: sending RST, seq 0, ack 3312452185
*Jul 20 13:19:00.099: TCP: sent RST to 192.168.1.2:25147 from 1.1.1.1:179
*Jul 20 13:19:00.099: Released port 0 in Transport Port Agent for TCP IP type 0 delay 240000
*Jul 20 13:19:00.099: TCP0: state was LISTEN -> CLOSED [0 -> UNKNOWN(0)]
*Jul 20 13:19:00.103: TCB 0x682D4F80 destroyed
```

COMMAND:

```
neighbos x.x.x.x update-source loopback x
```

upvoted 3 times

Almylle 5 months, 3 weeks ago

Selected Answer: D

Im agree with HungarianDish, the command update source in bgp, is important to have a TCP communication between iBGP neighbors, in this case in the debug they router is trying to open a tcp connection but it's failed because of the missing command from PE3

upvoted 3 times

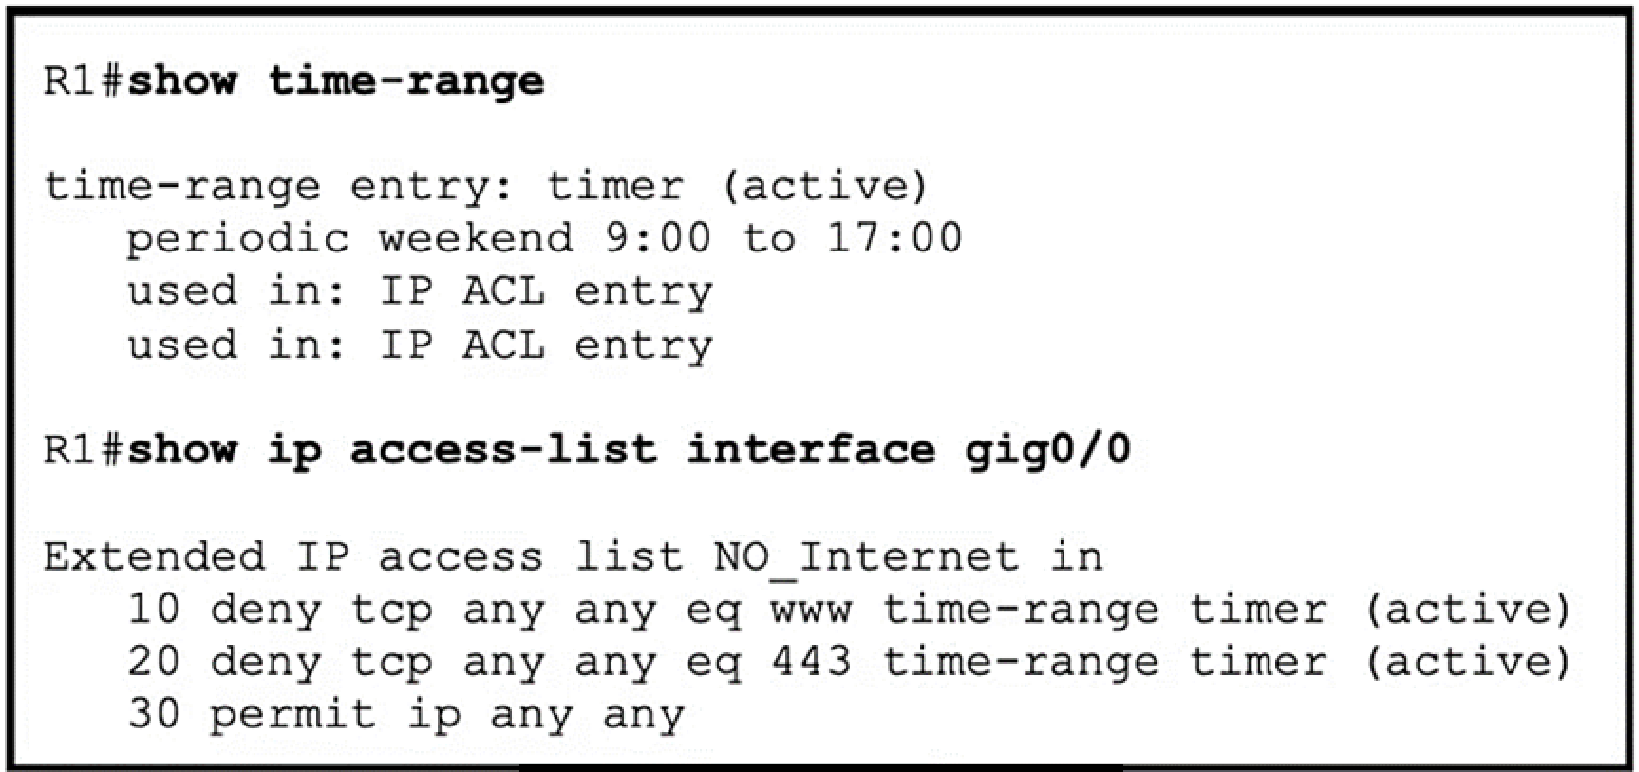

```
R1#show time-range

time-range entry: timer (active)
   periodic weekend 9:00 to 17:00
   used in: IP ACL entry
   used in: IP ACL entry

R1#show ip access-list interface gig0/0

Extended IP access list NO_Internet in
   10 deny tcp any any eq www time-range timer (active)
   20 deny tcp any any eq 443 time-range timer (active)
   30 permit ip any any
```

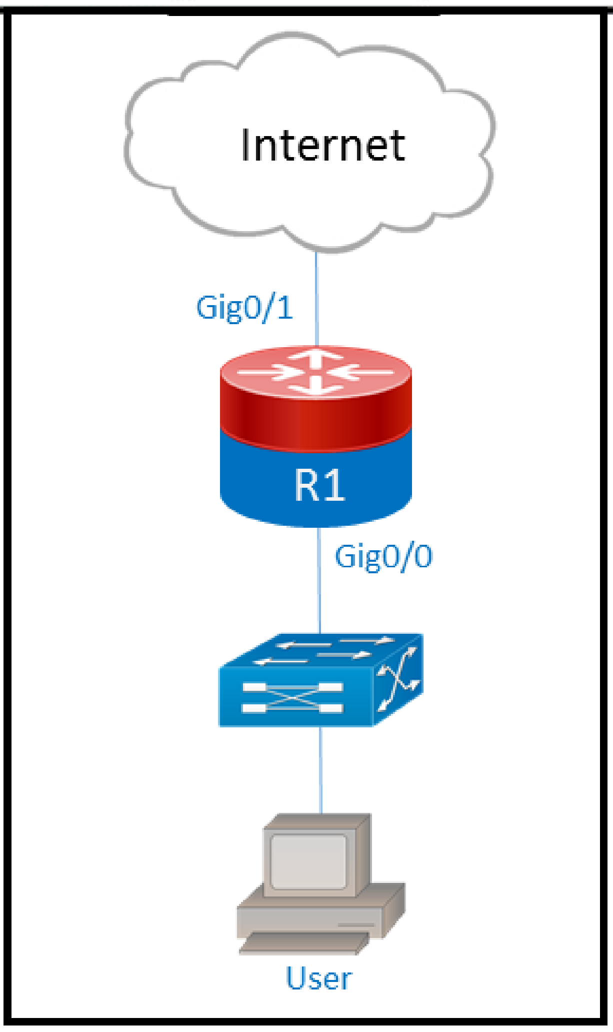

Refer to the exhibit. Users on a call center report that they cannot browse the internet on Saturdays during the afternoon. Which configuration resolves the issue?

A. time-range timer
no periodic weekend 9:00 to 17:00
periodic weekend 17:00 to 23:59

B. no time-range timer

C. interface gig0/0
ip access-group NO_Internet out

D. ip access-list extended NO_Internet
15 permit tcp any any eq www

Correct Answer: *A*

inteldarvid 5 months, 1 week ago

Selected Answer: A

its logical. A

upvoted 3 times

Opreis 3 months, 1 week ago

i dont think afternoon is between 1700 - 23:59

upvoted 2 times

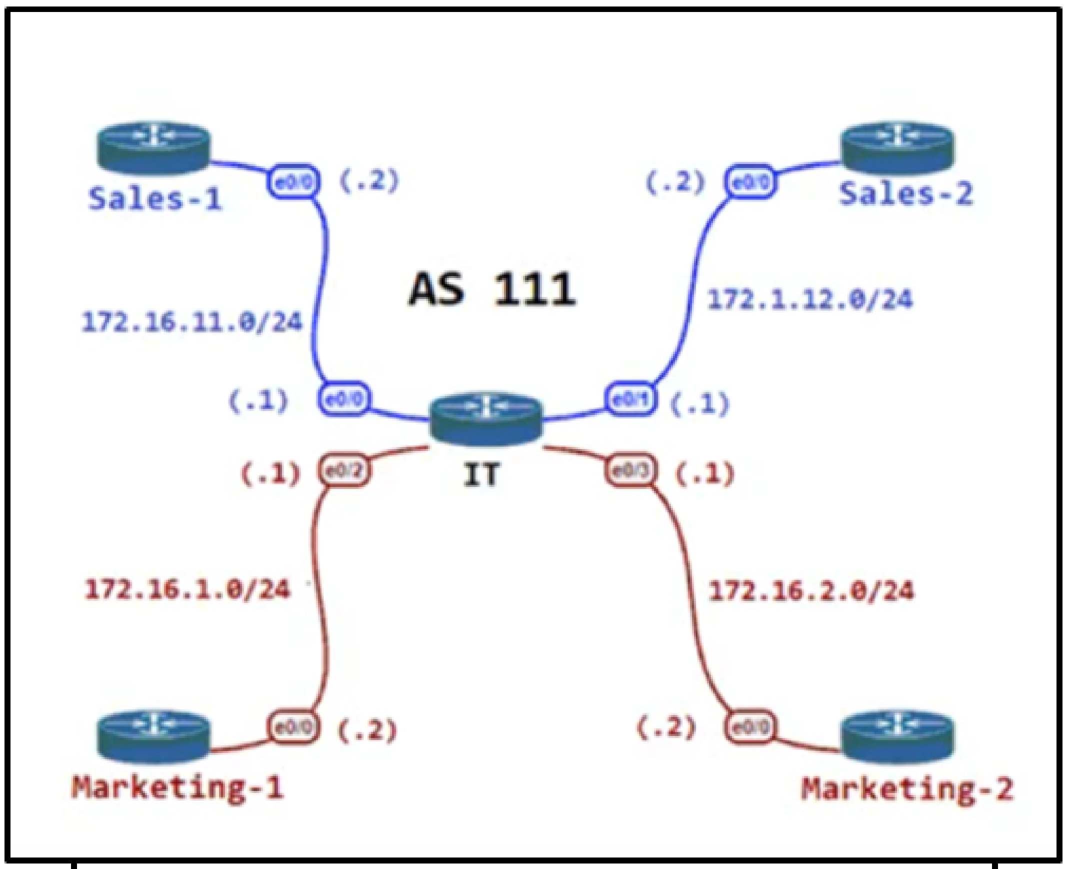

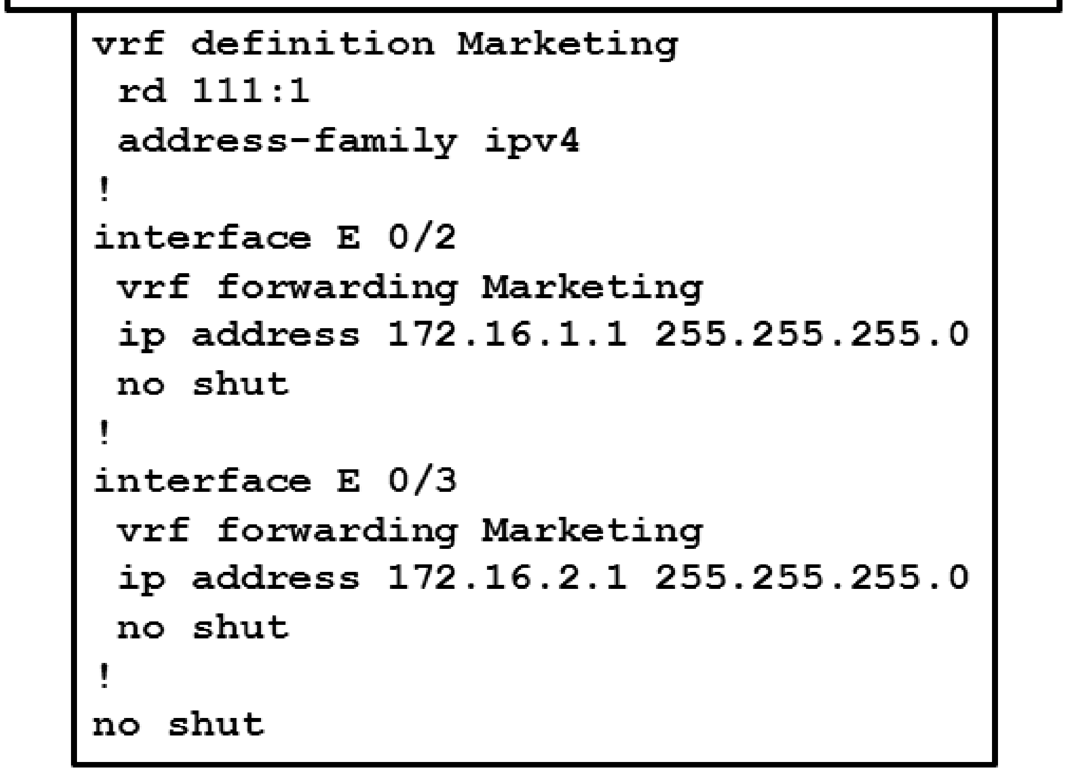

```
vrf definition Marketing
 rd 111:1
 address-family ipv4
!
interface E 0/2
 vrf forwarding Marketing
 ip address 172.16.1.1 255.255.255.0
 no shut
!
interface E 0/3
 vrf forwarding Marketing
 ip address 172.16.2.1 255.255.255.0
 no shut
!
no shut
```

Refer to the exhibit. The IT router is connected with the Sales and Marketing departments. The interfaces have been assigned to their respective VRFs to keep the two department routes isolated. Which configuration set must the IT router use for BGP to distribute routes for each department that maintains their own routing table for network isolation?

A. router bgp 111
address-family ipv4 unicast
neighbor 172.16.1.2 remote-as 111
neighbor 172.16.2.2 remote-as 111
neighbor 172.16.11.2 remote-as 111
neighbor 172.16.12.2 remote-as 111

B. router bgp 111
address-family ipv4 vrf Marketing
neighbor 172.16.1.2 remote-as 111
neighbor 172.16.2.2 remote-as 111
!
address-family ipv4 vrf Sales
neighbor 172.16.11.2 remote-as 111
neighbor 172.16.12.2 remote-as 111

```
C. router bgp 111
neighbor 172.16.1.2 remote-as 111
neighbor 172.16.2.2 remote-as 111
neighbor 172.16.11.2 remote-as 111
neighbor 172.16.12.2 remote-as 111
```

```
D. router bgp 111
address-family ipv4 vrf Marketing
neighbor 172.16.1.2 remote-as 111
neighbor 172.16.1.2 Route-reflector-client
neighbor 172.16.2.2 remote-as 111
neighbor 172.16.2.2 Route-reflector-client
!
address-family ipv4 vrf Sales
neighbor 172.16.11.2 remote-as 111
neighbor 172.16.11.2 Route-reflector-client
neighbor 172.16.12.2 remote-as 111
neighbor 172.16.12.2 Route-reflector-client
```

Correct Answer: *B*

potato_inet0 Highly Voted 👍 7 months ago

D is the correct answer, we have iBGP setup, Routes received by the IT router will be iBP routes so they'll never be sent to iBGP, we need route-reflection per each VRF.

upvoted 7 times

guy276465281819372 Most Recent ⏲ 4 months, 3 weeks ago

Selected Answer: D

D correct

upvoted 1 times

[Removed] 4 months, 3 weeks ago

Selected Answer: D

B Could have worked if it had "next-hop-self" for each of the neighbors.

D is the correct answer

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: D

yes option D correct, Because its necesary use router reflector in IBGP

upvoted 1 times

MEDO95 6 months, 2 weeks ago

I vote for D, B isn’t completed conf

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: D

I agree with potato_inet0. Also confirmed solution "D" in the lab. Connection is established between iBGP peers, but the routes advertised by the edge routers are only available if the central IT router is setting the peers as route-reflector-client. For instance, if Marketing-1 is advertising a LAN with ip 1.1.1.1/24, Marketing-2 can only reach that if they both are RRclients and IT is RR.

upvoted 3 times

HungarianDish 6 months, 3 weeks ago

The central router IT needs to be set as RR for full connectivity in each vrf. So, answer "D" is valid.

upvoted 2 times

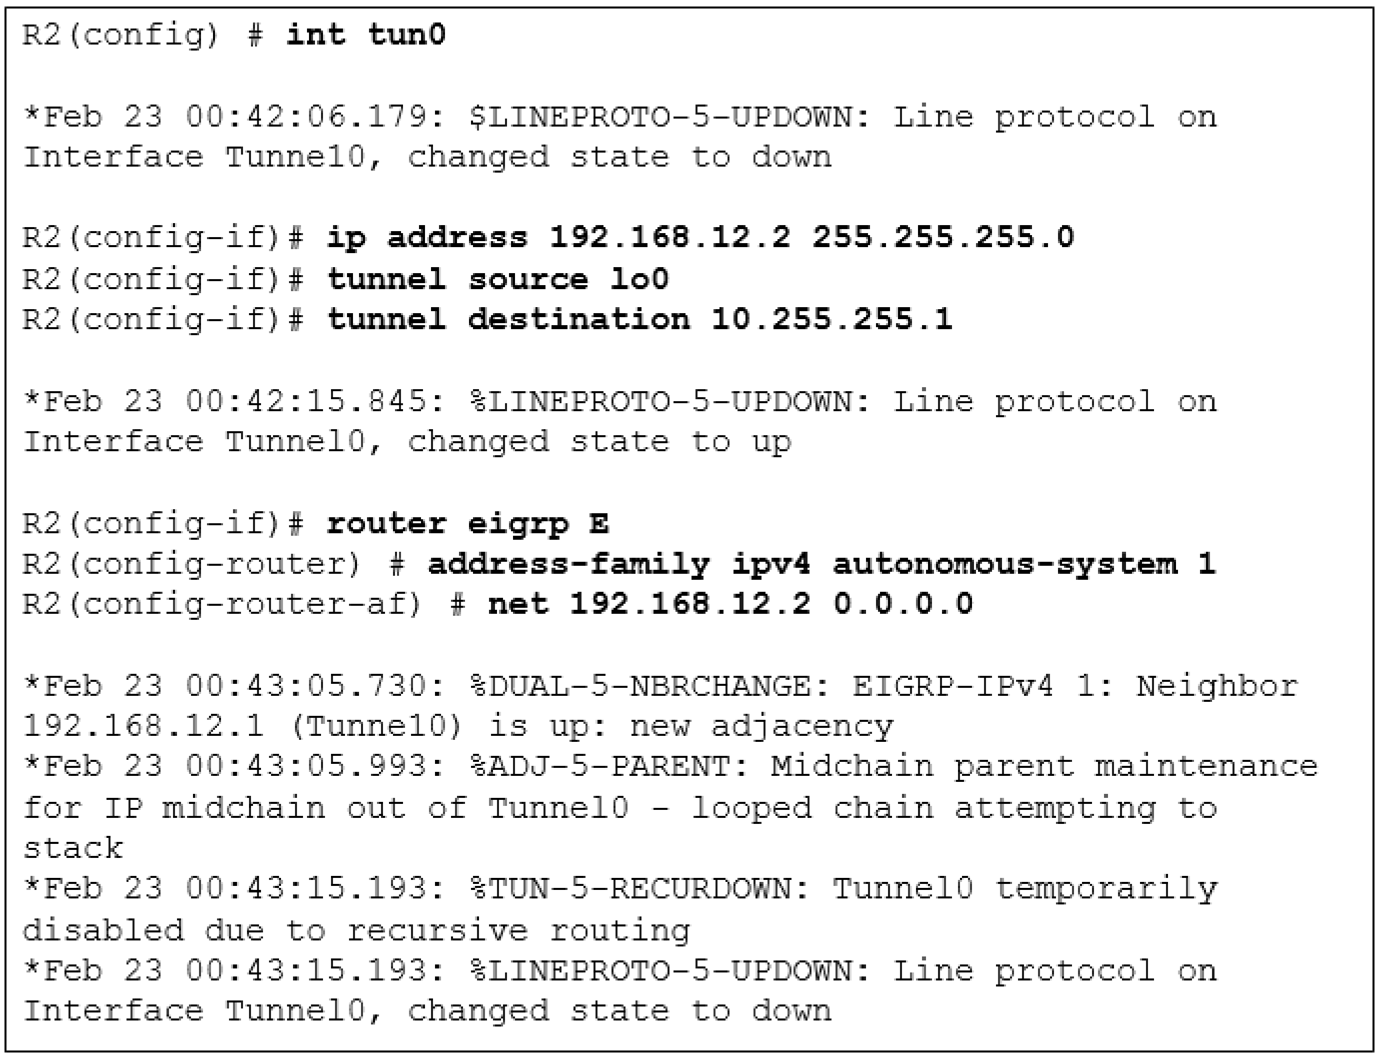

```
R2(config) # int tun0

*Feb 23 00:42:06.179: $LINEPROTO-5-UPDOWN: Line protocol on
Interface Tunnel0, changed state to down

R2(config-if)# ip address 192.168.12.2 255.255.255.0
R2(config-if)# tunnel source lo0
R2(config-if)# tunnel destination 10.255.255.1

*Feb 23 00:42:15.845: %LINEPROTO-5-UPDOWN: Line protocol on
Interface Tunnel0, changed state to up

R2(config-if)# router eigrp E
R2(config-router) # address-family ipv4 autonomous-system 1
R2(config-router-af) # net 192.168.12.2 0.0.0.0

*Feb 23 00:43:05.730: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor
192.168.12.1 (Tunnel0) is up: new adjacency
*Feb 23 00:43:05.993: %ADJ-5-PARENT: Midchain parent maintenance
for IP midchain out of Tunnel0 - looped chain attempting to
stack
*Feb 23 00:43:15.193: %TUN-5-RECURDOWN: Tunnel0 temporarily
disabled due to recursive routing
*Feb 23 00:43:15.193: %LINEPROTO-5-UPDOWN: Line protocol on
Interface Tunnel0, changed state to down
```

Refer to the exhibit. An administrator is configuring a GRE tunnel to establish an EIGRP neighbor to a remote router. The other tunnel endpoint is already configured. After applying the configuration as shown, the tunnel started flapping. Which action resolves the issue?

A. Modify the network command to use the Tunnel0 interface netmask.

B. Stop sending a route matching the tunnel destination across the tunnel.

C. Advertise the Loopback0 interface from R2 across the tunnel.

D. Readdress the IP network on the Tunnel0 on both routers using the /31 netmask.

Correct Answer: *B*

⊟ **inteldarvid** 4 months, 3 weeks ago

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/22327-gre-flap.html

upvoted 1 times

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: B

B correct:

In this question we are advertising the tunnel IP address 192.168.12.2 to the other side. When other end receives the EIGRP advertisement, it realizes it can reach the other side of the tunnel via EIGRP.
In other words, it reaches the tunnel destination through the tunnel itself -> This causes “recursive routing” error.
Note: In order to avoid this error, do not advertise the tunnel destination IP address on the tunnel interface to other side.

upvoted 2 times

⊟ **HungarianDish** 6 months, 2 weeks ago

The output is not completely relevant to the question. It only shows that the tunnel network is advertised, which is OK. It should not advertise 10.255.255.1 0.0.0.0 (the physical address/WAN address of the peering) under the same eigrp process, because that would cause a loop.

upvoted 3 times

Question #454 *Topic 1*

An engineer must override the normal routing behavior of a router. The engineer must send HTTP traffic that is destined to 10.100.100.100 from 10.1.1.0/24 via a next hop of 10.2.2.2, two hops away from the router that is connected to the 10.1.1.0/24 subnet. Which configuration reroutes traffic according to this requirement?

A.
```
access-list 100 permit tcp 10.1.1.0 0.0.0.255 host 10.100.100.100 eq http
!
route-map POLICY permit 10
match ip address 100
set ip next-hop recursive 10.2.2.2
```

B.
```
access-list 100 permit tcp 10.1.1.0 0.0.0.255 host 10.100.100.100 eq http
!
route-map POLICY deny 10
match ip address 100
set ip next-hop recursive 10.2.2.2
route-map POLICY permit 20
```

C.
```
access-list 100 permit tcp 10.1.1.0 0.0.0.255 host 10.100.100.100 eq http
!
route-map POLICY permit 10
match ip address 100
set ip next-hop 10.2.2.2
route-map POLICY permit 20
```

D.
```
access-list 100 permit tcp 10.1.1.0 0.0.0.255 host 10.100.100.100 eq http
!
route-map POLICY permit 10
match ip address 100
set ip next-hop 10.2.2.2
```

Correct Answer: *A*

inteldarvid 5 months, 1 week ago

Selected Answer: A

yes, correct "A". But also its necesary permit route map seq 20. :)

upvoted 2 times

> **DeWalt95** 5 days, 14 hours ago
>
> I believe route maps for PBR do not require a permit statement
>
> upvoted 1 times
>
> **[Removed]** 4 months, 3 weeks ago
>
> Correct, otherwise all other traffic is blackholed
>
> upvoted 1 times
>
> **Brand** 3 months, 3 weeks ago
>
> What you're describing is the answer B. Is it B then?
>
> upvoted 1 times
>
> > **Brand** 3 months, 3 weeks ago
> >
> > Sorry, B is denying the HTTP traffic in Seq 10. It should be "A"
> >
> > upvoted 1 times

Question #455 *Topic 1*

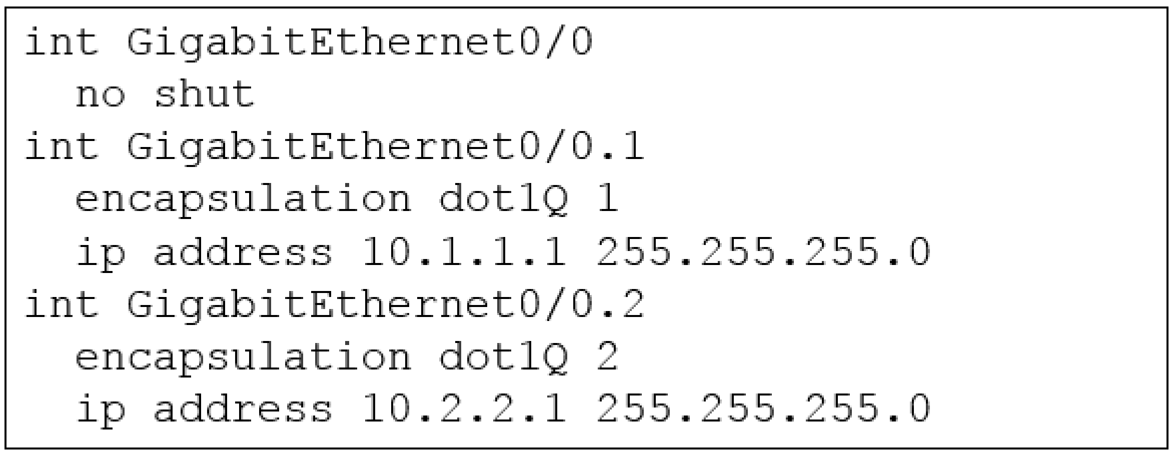

```
int GigabitEthernet0/0
  no shut
int GigabitEthernet0/0.1
  encapsulation dot1Q 1
  ip address 10.1.1.1 255.255.255.0
int GigabitEthernet0/0.2
  encapsulation dot1Q 2
  ip address 10.2.2.1 255.255.255.0
```

Refer to the exhibit. Two routers are connected back to back via Gigabit Ethernet 0/0 interfaces. Which configuration provides VRF-Lite connectivity for two separate VRFs using the prefixes 10.1.1.0/24 for one VRF and 10.2.2.0/24 for the other VRF?

A. ip vrf 1
rd 65001:1
ip vrf 2
rd 65001:2
!
int GigabitEthernet0/0
no shut
!
int GigabitEthernet0/0.1
encapsulation dot1Q 1
ip vrf forwarding
ip address 10.1.1.1 255.255.255.0
!
int GigabitEthernet0/0.2
encapsulation dot1Q 2
ip vrf forwarding
ip address 10.2.2.1 255.255.255.0

B. ip vrf 1
rd 65001:1
ip vrf 2
rd 65001:1
!
int GigabitEthernet0/0
no shut
!
int GigabitEthernet0/0.1
encapsulation dot1Q 1
ip vrf forwarding 1
ip address 10.1.1.1 255.255.255.0
!
int GigabitEthernet0/0.2
encapsulation dot1Q 2
ip vrf forwarding 2
ip address 10.2.2.1 255.255.255.0

C. ip vrf 1
ip vrf 2
int GigabitEthernet0/0
no shut
!
int GigabitEthernet0/0.1
encapsulation dot1Q 1

```
ip vrf forwarding 1
ip address 10.1.1.1 255.255.255.0
!
int GigabitEthernet0/0.2
encapsulation dot1Q 2
ip vrf forwarding 2
ip address 10.2.2.1 255.255.255.0
```

```
D. ip vrf 1
ip vrf 2
!
int GigabitEthernet0/0
no shut
!
int GigabitEthernet0/0.1
encapsulation dot1Q 1
ip address 10.1.1.1 255.255.255.0
ip vrf forwarding 1
!
int GigabitEthernet0/0.2
encapsulation dot1Q 2
ip address 10.2.2.1 255.255.255.0
ip vrf forwarding 2
```

Correct Answer: *C*

aqwsdfghjklp 2 weeks, 1 day ago

Why not B?
rd should be mandatory.

upvoted 1 times

Brand 3 months, 2 weeks ago

Selected Answer: C

Just wanted to say that the option A is literally impossible as the router not going to you set "ip vrf forwarding" under an interface without defining the vrf itself. So it's IMPOSIBBLE to see such an output. Try and you'll see.

upvoted 1 times

guy276465281819372 4 months, 3 weeks ago

Selected Answer: C

C correct

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

yes option ""C"" correct, beacuse ir vrf lite without MP-BGP

upvoted 1 times

sajjad_gayyem 5 months, 3 weeks ago

Selected Answer: A

RD definition is preferred.

upvoted 1 times

> **inteldarvid** 5 months, 1 week ago
>
> you are wrong my friend, because. you need name vrf: ip vrf forwarding "1" or 2
>
> upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: C

https://zartmann.dk/mpls-vpns-vs-vrf-lite/

upvoted 3 times

> **HungarianDish** 6 months, 2 weeks ago
>
> Some recent IOS requires the RD to be set also for VRF-Lite. It is only locally significant though.
> https://www.packetcoders.io/cisco-ios-how-to-configure-vrf-lite/
>
> upvoted 1 times

Question #456 Topic 1

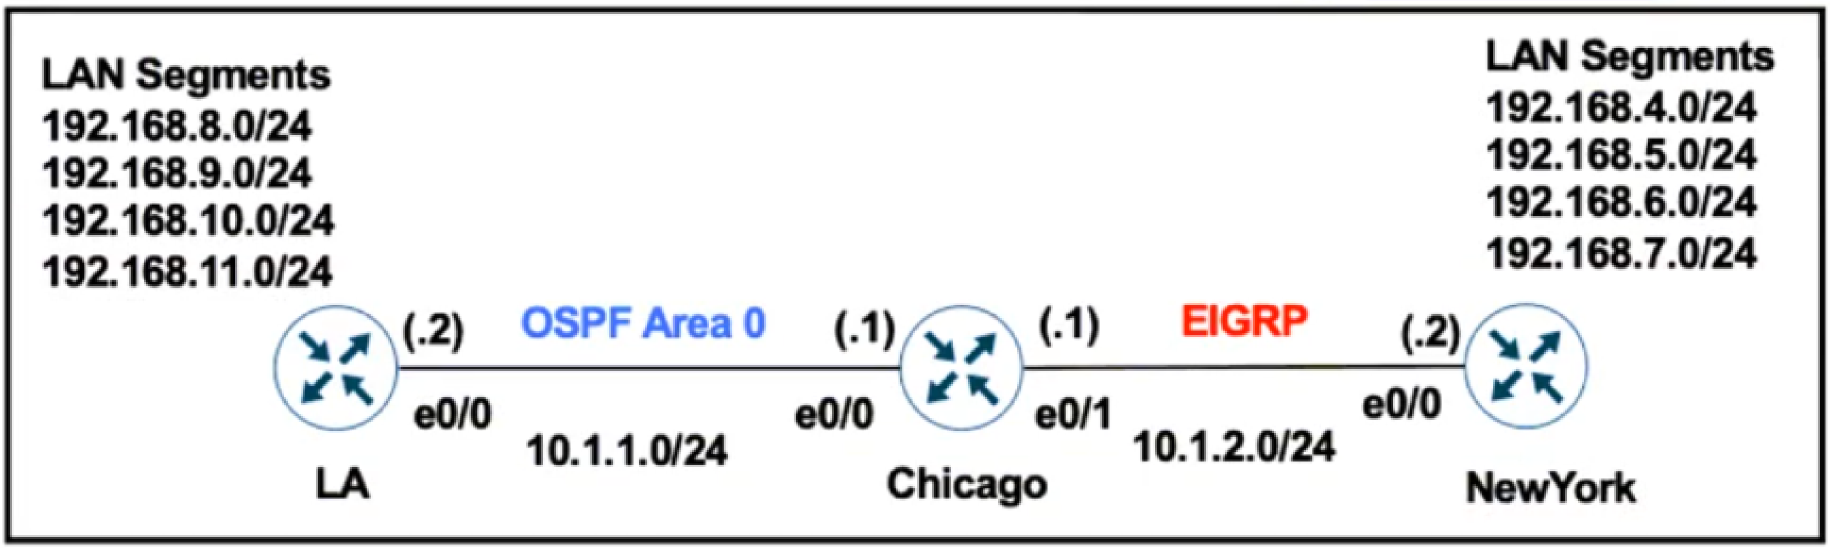

Refer to the exhibit. The network administrator configures the Chicago router to mutually redistribute the LA and New York routes with EIGRP routes to be summarized as a single route in OSPF:

```
router eigrp 100
redistribute ospf 1 metric 10 10 10 10 10
router ospf 1
redistribute eigrp 100 subnets
area 0 range 192.168.4.0 255. 255.252.0
```

After the configuration, the LA router receives all the specific New York routes except the summary route. Which set of configurations resolve the issue on the Chicago router?

A. interface E 0/0
ip summary-address eigrp 100 192.168.4.0 255.255.252.0

B. router ospf 1
summary-address 192.168.4.0 255.255.252.0

C. interface E 0/0
summary-address 192.168.4.0 255.255.252.0

D. router ospf 1
area 0 range 192.168.0.0 255.255.0.0

Correct Answer: *B*

[Removed] 4 months, 2 weeks ago

Selected Answer: B

Remember that ABRs benefit from the command "area <area-id> range <summary>
This command means that from this Area, advertise this range.
But Chicago Router is an ASBR, which requires the command "summary-address <summary-prefix>" to summarize routes into the OSPF domain.

upvoted 1 times

[Removed] 4 months, 3 weeks ago

Selected Answer: B

A. Only for EIGRP links, E0/0 is an OSPF link
B. ASBRs require "summary-address <summary-address>" to summarize external routes.

upvoted 2 times

inteldarvid 5 months, 1 week ago

Selected Answer: B

yes, option "B" correct

upvoted 1 times

HungarianDish 7 months, 1 week ago

Selected Answer: B

https://www.geeksforgeeks.org/configuring-ospf-route-summarization-in-cisco/
ASBR summary-address <network-id> <prefix-mask> [not-advertise]
ABR area <area-id> range <network-id> <prefix-mask> [advertise |not-advertise]

upvoted 3 times

Question #457 *Topic 1*

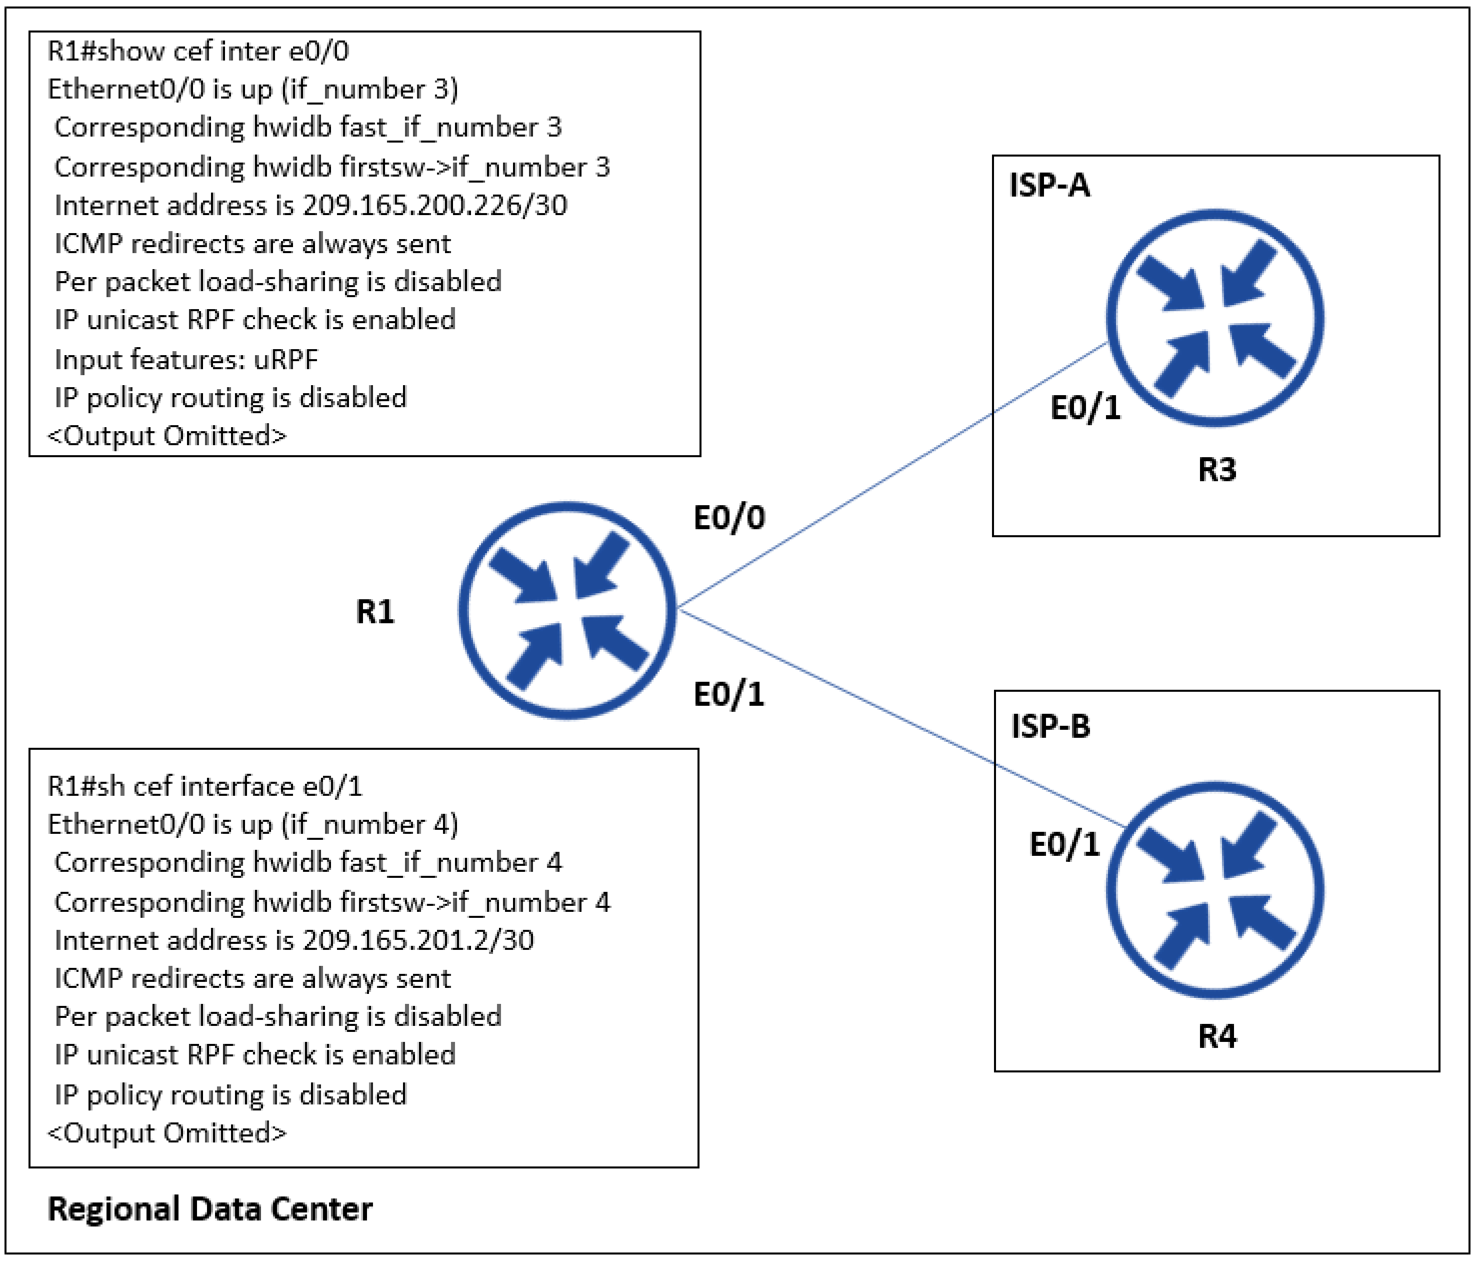

```
R1#show cef inter e0/0
Ethernet0/0 is up (if_number 3)
 Corresponding hwidb fast_if_number 3
 Corresponding hwidb firstsw->if_number 3
 Internet address is 209.165.200.226/30
 ICMP redirects are always sent
 Per packet load-sharing is disabled
 IP unicast RPF check is enabled
 Input features: uRPF
 IP policy routing is disabled
<Output Omitted>
```

```
R1#sh cef interface e0/1
Ethernet0/0 is up (if_number 4)
 Corresponding hwidb fast_if_number 4
 Corresponding hwidb firstsw->if_number 4
 Internet address is 209.165.201.2/30
 ICMP redirects are always sent
 Per packet load-sharing is disabled
 IP unicast RPF check is enabled
 IP policy routing is disabled
<Output Omitted>
```

Refer to the exhibit. The company implemented uRPF to address an antispoofing attack. A network engineer received a call from the IT security department that the regional data center is under an IP spoofing attack. Which configuration must be implemented on R1 to resolve this issue?

A. interface ethernet0/0
ip verify unicast reverse-path

B. interface ethernet0/1
ip verify unicast reverse-path

C. interface ethernet0/0
ip unicast RPF check reachable-via any allow-default allow-self-ping

D. interface ethernet0/1
ip unicast RPF check reachable-via any allow-default allow-self-ping

Correct Answer: *B*

inteldarvid 5 months, 1 week ago

Selected Answer: B

o9ption corerct is B, is easy beacuse: the command is:
#ip verify unicast source reachable-via any allow-self-ping (its works), and I can see in 0/0 interface 0/0 "input features: uRPF". I have only one option: Option "B"

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

As uRPF is already enabled on both interfaces (see "ip unicast RPF check is enabled" under "show cef int"), it is hard to choose between "A" and "B".

upvoted 2 times

HungarianDish 7 months ago

This is a vague question. uRPF is already enabled for both WAN interfaces, as shown in the output under show cef int ... (ip unicast RPF check is enabled)
"ip verify unicast reverse-path" is the old command for strict mode. The new commands are recommended.
Plus, it is a multihome environemnt, where loose mode would be appropriate instead of strict.

"C","D": ip unicast RPF check reachable-via any allow-default allow-self-ping => allow-default + loose mode makes no sense on the internet facing interfaces.
ip unicast RPF check reachable-via any allow-self-ping => Loose mode allowing local device to ping it's own interface would be OK, but it's not an option.
Probably, in real exam we can choose both "A" and "B".

upvoted 1 times

HungarianDish 7 months ago

ip unicast RPF check reachable-via any allow-default allow-self-ping
https://learningnetwork.cisco.com/s/question/0D53i00000Kt5tDCAR/urpf-allowdefault

Using loose mode with allow-default can in some (if not most) cases completely defeat the purpose of implementing uRPF at all.

https://community.cisco.com/t5/other-security-subjects/urpf-with-default-route/td-p/1182324

having a default-route will pretty much negate the use of uRPF as the router will always have a path back to the source.

upvoted 1 times

HungarianDish 7 months ago

https://www.cisco.com/c/en/us/td/docs/routers/asr920/configuration/guide/sec-data-urpf/17-1-1/b-sec-data-urpf-xe-17-1-asr920/b-sec-data-urpf-xe-17-1-asr920_chapter_01.pdf
Loose Mode for dual-homed ISP
-alleviates the interface dependency of strict mode

upvoted 1 times

Question #458 Topic 1

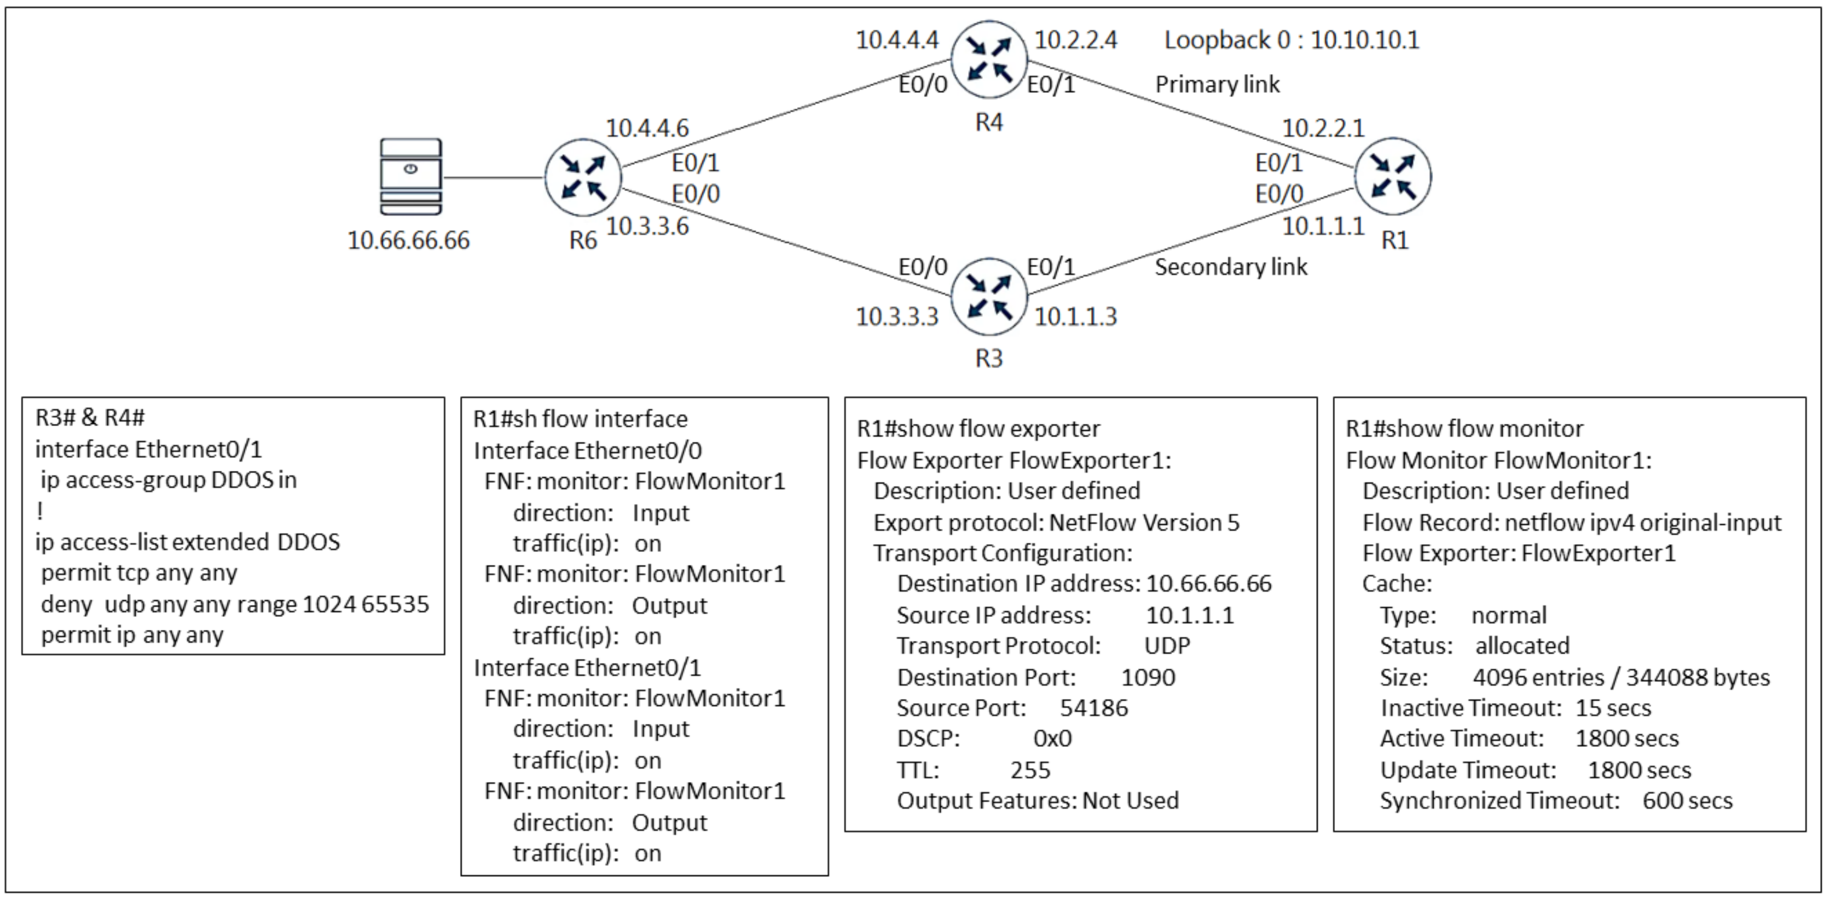

```
R3# & R4#
interface Ethernet0/1
 ip access-group DDOS in
!
ip access-list extended DDOS
 permit tcp any any
 deny  udp any any range 1024 65535
 permit ip any any
```

```
R1#sh flow interface
Interface Ethernet0/0
 FNF: monitor: FlowMonitor1
      direction:   Input
      traffic(ip):  on
 FNF: monitor: FlowMonitor1
      direction:   Output
      traffic(ip):  on
Interface Ethernet0/1
 FNF: monitor: FlowMonitor1
      direction:   Input
      traffic(ip):  on
 FNF: monitor: FlowMonitor1
      direction:   Output
      traffic(ip):  on
```

```
R1#show flow exporter
Flow Exporter FlowExporter1:
  Description: User defined
  Export protocol: NetFlow Version 5
  Transport Configuration:
    Destination IP address: 10.66.66.66
    Source IP address:        10.1.1.1
    Transport Protocol:     UDP
    Destination Port:       1090
    Source Port:      54186
    DSCP:            0x0
    TTL:          255
    Output Features: Not Used
```

```
R1#show flow monitor
Flow Monitor FlowMonitor1:
  Description: User defined
  Flow Record: netflow ipv4 original-input
  Flow Exporter: FlowExporter1
  Cache:
    Type:      normal
    Status:   allocated
    Size:      4096 entries / 344088 bytes
    Inactive Timeout:  15 secs
    Active Timeout:     1800 secs
    Update Timeout:     1800 secs
    Synchronized Timeout:    600 secs
```

Refer to the exhibit. An engineer configured NetFlow but cannot receive the flows from R1. Which two configurations resolve the issue? (Choose two.)

A. R3(config)#ip access-list extended DDOS
R3(config-ext-nacl)#5 permit udp any host 10.66.66.66 eq 1090

B. R4(config)#flow exporter FlowExporter1
R4(config-flow-exporter)#destination 10.66.66.66

C. R3(config)#flow exporter FlowExporter1
R3(config-flow-exporter)#destination 10.66.66.66

D. R1(config)#flow exporter FlowExporter1
R1(config-flow-exporter)#destination 10.66.60.66

E. R4(config)#ip access-list extended DDOS
R4(config-ext-naci)#5 permit udp any host 10.66.66.66 eq 1090

Correct Answer: *AE*

aqwsdfghjklp 5 days, 15 hours ago
In fact, isn't A enough?
Tip:source IP address 10.1.1.1
upvoted 1 times

chris110 3 months, 1 week ago
WATCH OUT! Maybe the exhibition is not accurate! Saw it diffrent on networktut.
upvoted 1 times

inteldarvid 5 months, 1 week ago
correct
upvoted 1 times

Almylle 5 months, 3 weeks ago
The given answer is correct, because the port of othe destionation from the router is dropped because of the access list on R3 and R4, so u have to modify the ACL to permit that port.
upvoted 1 times

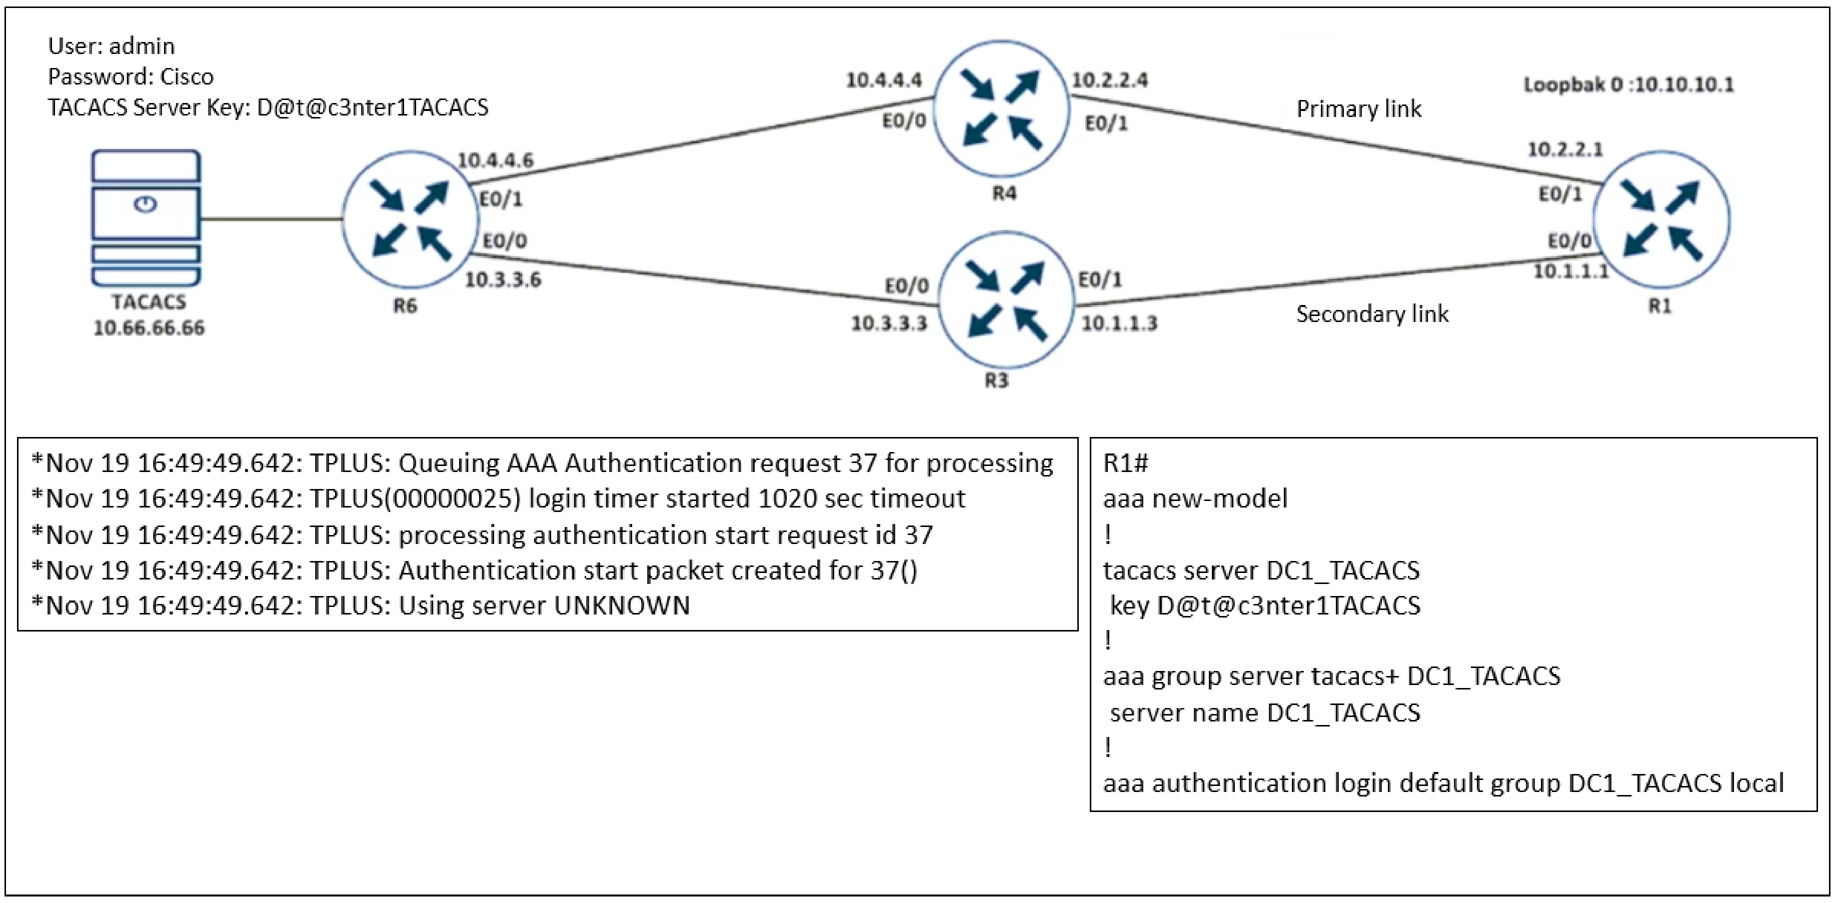

User: admin
Password: Cisco
TACACS Server Key: D@t@c3nter1TACACS

```
*Nov 19 16:49:49.642: TPLUS: Queuing AAA Authentication request 37 for processing
*Nov 19 16:49:49.642: TPLUS(00000025) login timer started 1020 sec timeout
*Nov 19 16:49:49.642: TPLUS: processing authentication start request id 37
*Nov 19 16:49:49.642: TPLUS: Authentication start packet created for 37()
*Nov 19 16:49:49.642: TPLUS: Using server UNKNOWN
```

```
R1#
aaa new-model
!
tacacs server DC1_TACACS
 key D@t@c3nter1TACACS
!
aaa group server tacacs+ DC1_TACACS
 server name DC1_TACACS
!
aaa authentication login default group DC1_TACACS local
```

Refer to the exhibit. R1 cannot authenticate via TACACS. Which configuration resolves the issue?

A. aaa group server tacacs+ DC1_TACACS
server name DC_TACACS

B. tacacs server DC1_TACACS
address ipv4 10.66.66.66
key D@t@c3nter1TACACS

C. tacacs server DC1_TACACS
address ipv4 10.60.66.66
key D@t@c3nter1TACACS

D. aaa group server tacacs+ DC_TACACS
server name DC_TACACS

Correct Answer: *B*

DeWalt95 2 weeks, 6 days ago

Selected Answer: B

B - its a spot the difference challenge.

upvoted 2 times

guy276465281819372 4 months, 3 weeks ago

am I blind or B AND C are the same answer?

upvoted 1 times

HarwinderSekhon 4 months ago

B has wrong IP of server

upvoted 2 times

[Removed] 4 months, 3 weeks ago

Spot the difference is always fun...

upvoted 3 times

inteldarvid 5 months, 1 week ago

Selected Answer: B

CORRECT

upvoted 1 times

Almylle 5 months, 3 weeks ago

The given answer is correct, u have to specify the server TACACS

upvoted 1 times

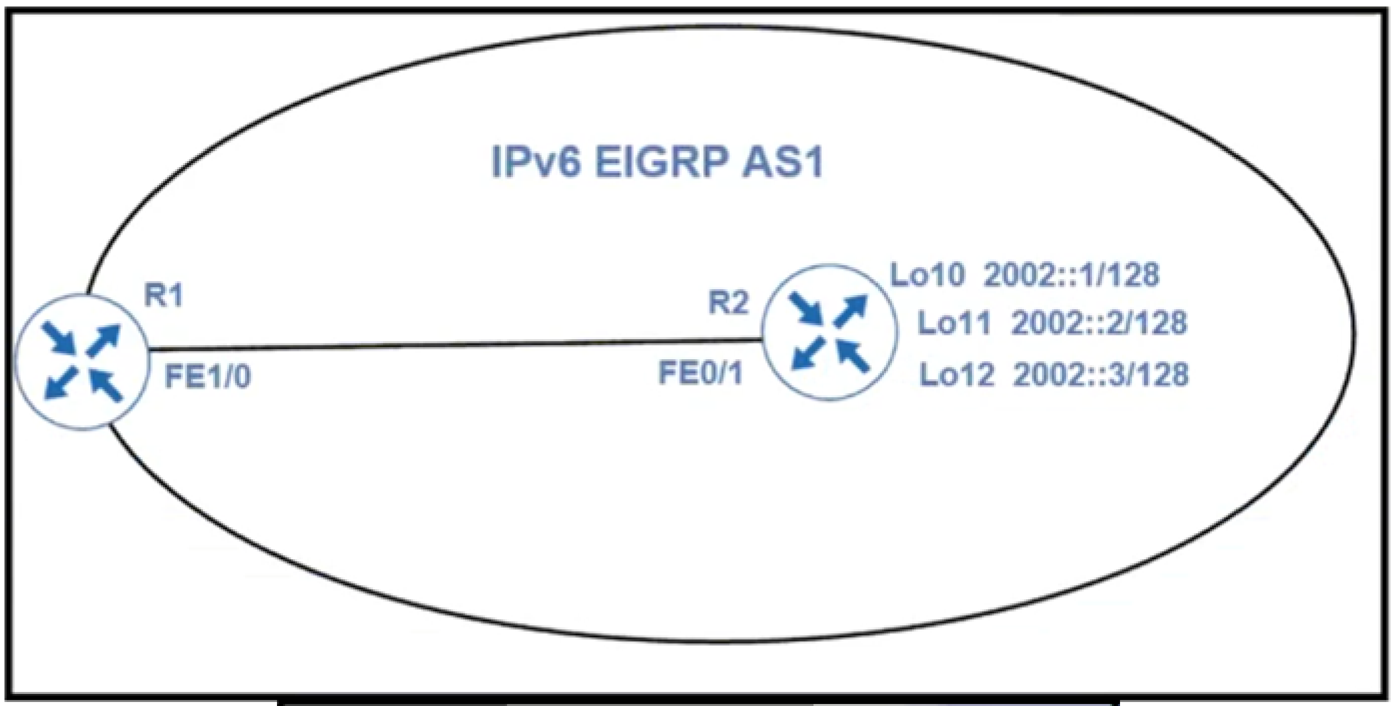

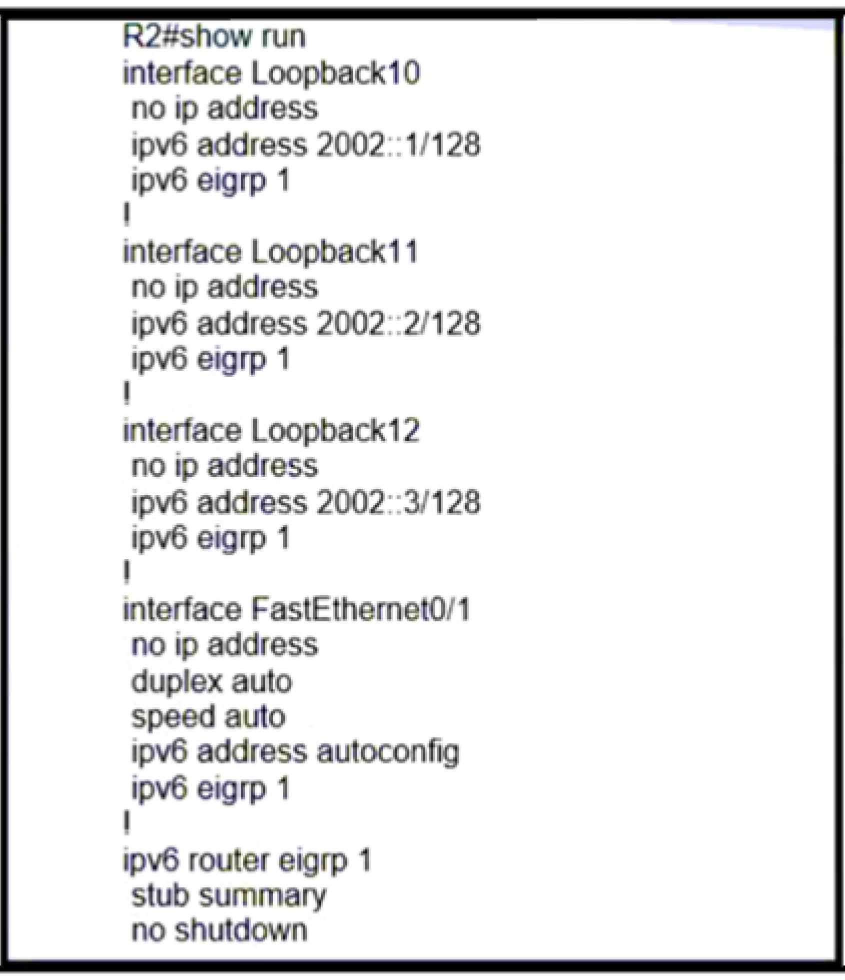

```
R2#show run
interface Loopback10
 no ip address
 ipv6 address 2002::1/128
 ipv6 eigrp 1
!
interface Loopback11
 no ip address
 ipv6 address 2002::2/128
 ipv6 eigrp 1
!
interface Loopback12
 no ip address
 ipv6 address 2002::3/128
 ipv6 eigrp 1
!
interface FastEthernet0/1
 no ip address
 duplex auto
 speed auto
 ipv6 address autoconfig
 ipv6 eigrp 1
!
ipv6 router eigrp 1
 stub summary
 no shutdown
```

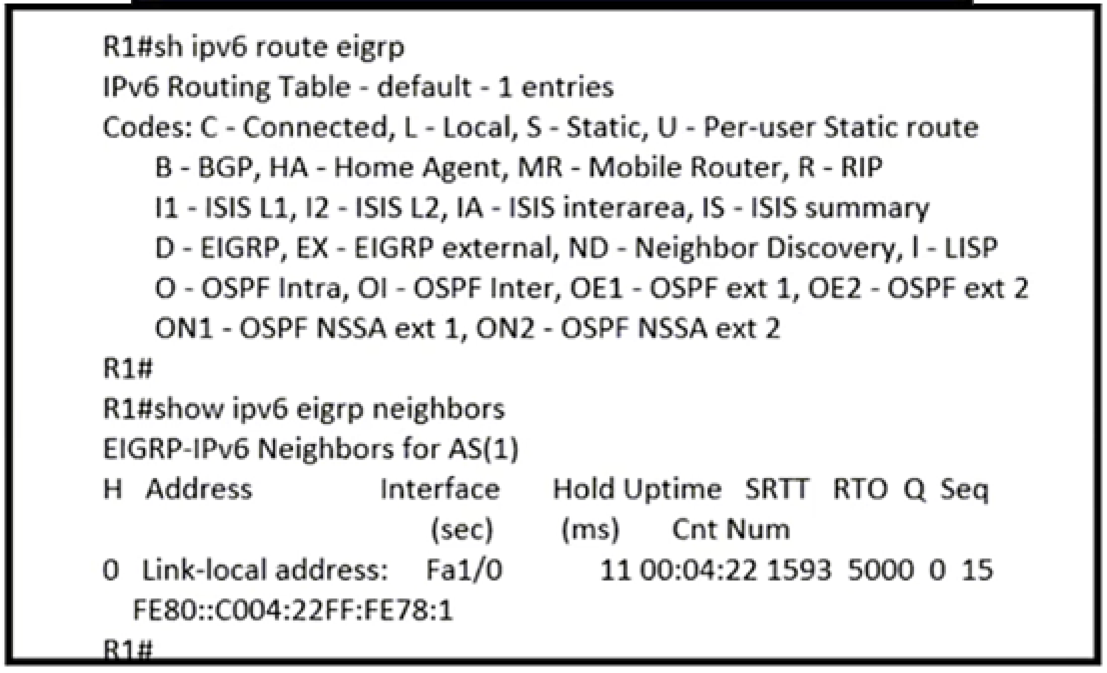

```
R1#sh ipv6 route eigrp
IPv6 Routing Table - default - 1 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
    B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
    I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
    D - EIGRP, EX - EIGRP external, ND - Neighbor Discovery, l - LISP
    O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
    ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
R1#
R1#show ipv6 eigrp neighbors
EIGRP-IPv6 Neighbors for AS(1)
H  Address           Interface      Hold Uptime   SRTT   RTO  Q  Seq
                                    (sec)         (ms)        Cnt Num
0  Link-local address:   Fa1/0        11 00:04:22 1593  5000  0  15
   FE80::C004:22FF:FE78:1
R1#
```

Refer to the exhibit. R1 cannot receive the R2 interfaces with individual prefixes. What must be reconfigured to advertise R2 interfaces to R1?

A. EIGRP process on R2 with the command stub summary receive-only

B. EIGRP process on R2 with the command stub summary connected

C. EIGRP process on R2 by removing the stub command keyword summary

D. interface FastEthernet0/1 on R2 with an EIGRP summary for all three loopback prefixes

Correct Answer: *B*

⊟ **aqwsdfghjklp** 1 month ago

The correct answer should be D, since the loopback address must also be redistributed

upvoted 1 times

⊟ **Muste** 4 months, 1 week ago

Selected Answer: B

B is the best answer

upvoted 1 times

⊟ **782f5f0** 4 months, 2 weeks ago

C

# eigrp stub ?
Connected - This means that the router can advertise directly connected networks that have been included into the EIGRP process

Leak-Map - A leak map allows you to fully customise which prefixes can be advertised to neighbours. Typically this is used in designs where a stub router is in transit for some traffic.

Receive-Only - This means that the router has no additional networks it knows about and should only receive updates. This is typically used on IDS appliances that connect to a single port on a switch only.

Redistributed - Allows the advertisement of external EIGRP prefixes into the domain.

Static - Allows the advertisement of static routes into the network.

Summary - Allows a summary route to be propagated across the network.

Carriage Return - By pressing ENTER this will accept the default of connected and summary only.

upvoted 1 times

⊟ **[Removed]** 4 months, 3 weeks ago

Selected Answer: C

Like Potato_inet0 commented. Both B and C work.
However, contrary to what Potato_inet0 is saying, the option C is not asking to remove the entire command "eigrp stub summary", it only states that it will remove only the keyword "summary". Option B is also correct, by adding the keyword "connected" the local routes to R-2 will be advertised.
Which option is best, I am not sure, as both options would cause EIGRP to reconverge and therefore a momentary outage.

upvoted 1 times

⊟ **inteldarvid** 5 months, 1 week ago

Selected Answer: B

option B, the connected keyword as it will advertise the loopbacks .

upvoted 2 times

⊟ **potato_inet0** 7 months ago

In this scenario both C and B are correct answers, however we already have a stub config in place so I think it's best not to delete the stub and just add the connected keyword as it will advertise the loopbacks .

upvoted 4 times

Question #461 *Topic 1*

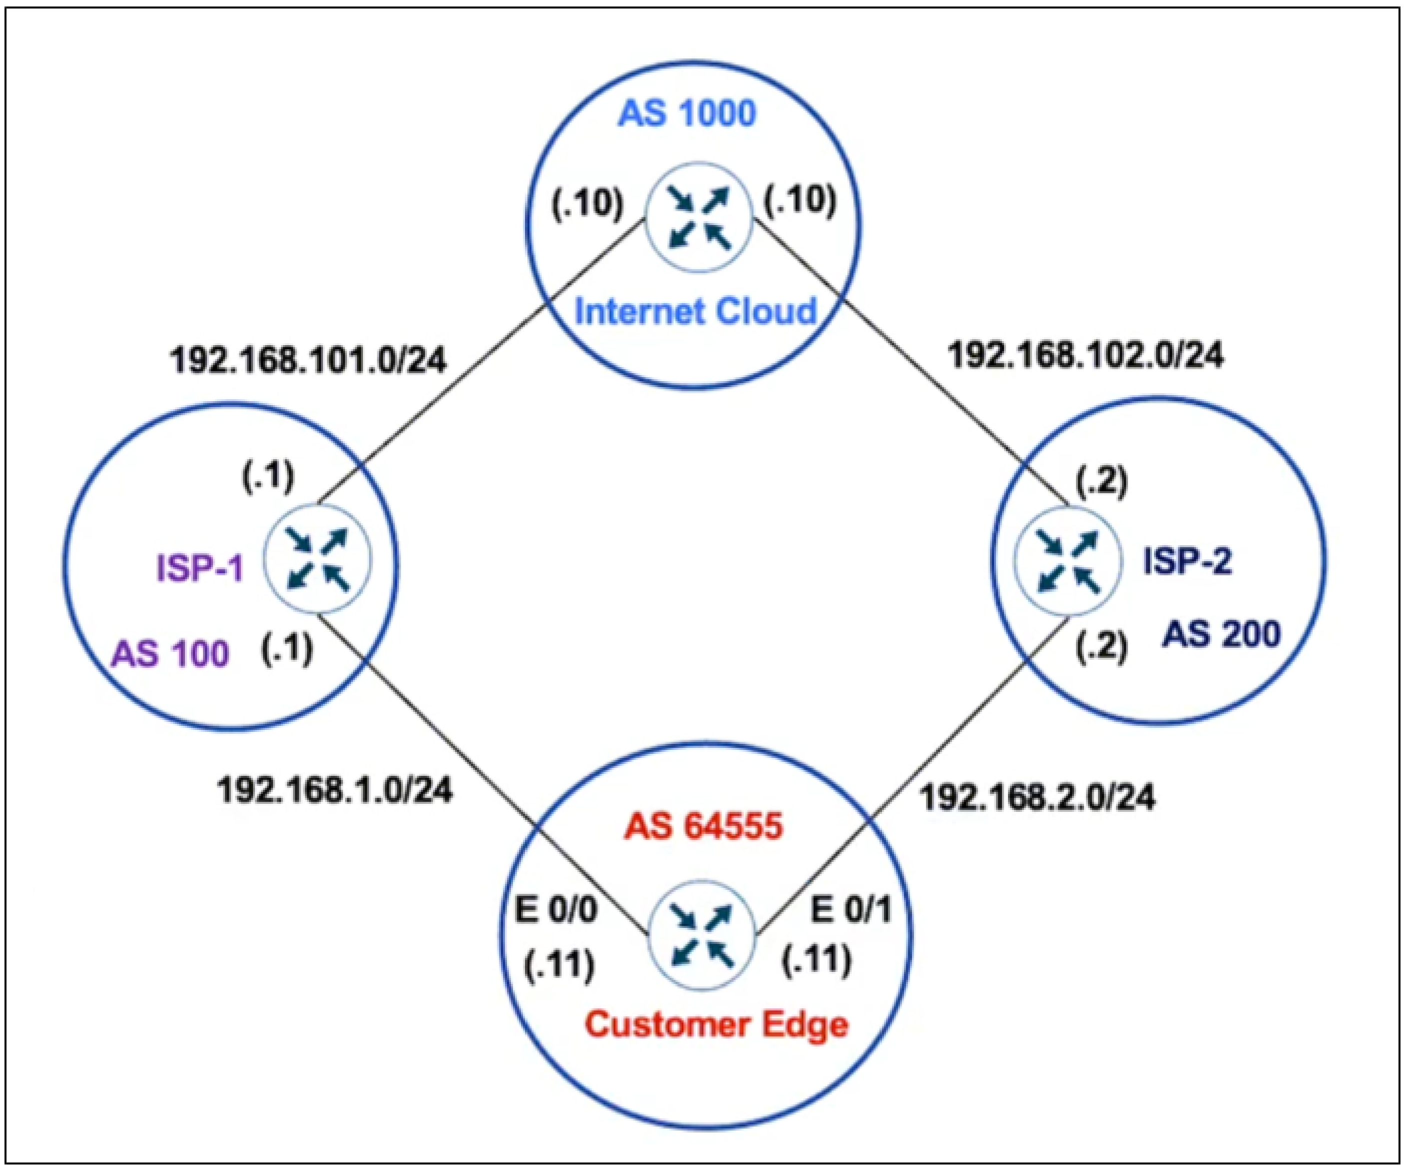

Refer to the exhibit. The Customer Edge router wants to use AS 100 as the preferred ISP for all external routes and SP2 as a backup.

Customer-Edge -

```
route-map SETAS
set as-path prepend 111
!
router bgp 64555
neighbor 192.168.1.1 remote-as 100
neighbor 192.168.2.2 remote-as 200
neighbor 192.168.2.2 route-map SETAS in
```

After this configuration, all the backup routes have disappeared from the BGP table on the Customer Edge router. Which set of configurations resolves the issue on the Customer Edge router?

A.
```
route-map SETAS
set as-path prepend 200
!
router bgp 64555
neighbor 192.168.1.1 remote-as 100
neighbor 192.168.2.2 remote-as 200
neighbor 192.168.2.2 route-map SETAS in
```

B.
```
route-map SETAS
set as-path prepend 111
!
router bgp 64555
neighbor 192.168.1.1 remote-as 100
neighbor 192.168.2.2 remote-as 200
neighbor 192.168.2.2 route-map SETAS out
```

```
C. route-map SETAS
set as-path prepend 200
!
router bgp 64555
neighbor 192.168.1.1 remote-as 100
neighbor 192.168.2.2 remote-as 200
neighbor 192.168.2.2 route-map SETAS out
```

```
D. route-map SETAS
set as-path prepend 111
!
router bgp 64555
neighbor 192.168.2.2 remote-as 100
neighbor 192.168.1.1 remote-as 200
neighbor 192.168.1.1 route-map SETAS in
```

Correct Answer: *A*

asans 1 week, 4 days ago

Something wrong with this question. The original route map config is the same as answer A, save for the different AS. And both works without issues

upvoted 1 times

yefrimart 2 months, 1 week ago

Selected Answer: A

When prepending an AS to a AS-PATH, you need to either prepend your own AS or the neighbor AS. Any other AS may cause problems.

For this question, because we want to influence outbound traffic, we need to prepend received prefixes. The answer that fulfill both conditions is answer A.

upvoted 1 times

spada05 4 months, 4 weeks ago

Selected Answer: A

Inbound, as the CE we presumably cannot touch the ISP and so have to manipulate inbound.

upvoted 2 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

for me option corerct is "C". I need out Prepending wit 200 200

upvoted 2 times

HungarianDish 6 months, 3 weeks ago

The original configuration (set as-path prepend 111) did not influence the path-selection at all. Plus, none of the routes disappeared. The question seems to be incorrect.

upvoted 2 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: A

Inbound is correct for sure. Only answer "A" uses correct neighbor + direction.
neighbor 192.168.2.2 route-map AS in

Probably this explains, why they want to prepend AS 200 (remote AS number):
https://blog.ipspace.net/2009/03/as-path-prepending-technical-details.html

Inbound AS-Path Prepending
... to add an extra copy of the REMOTE AS number to inbound updates received from (neighbor)

upvoted 3 times

HungarianDish 6 months, 3 weeks ago

Prepending as-number 111 inbound worked. It needs to be added two times, though (111 111).

```
CE#sh run | sec route-map
neighbor 192.168.2.2 route-map AS in
route-map AS permit 10
set as-path prepend 111 111
CE#
```

upvoted 1 times

> **HungarianDish** 6 months, 3 weeks ago
>
> This had the same effect:
> set as-path prepend 200 200

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Sorry, did not wait enough. Now it's OK. Using as-path prepend 111 111. Result: CE is choosing ISP1 as primary.

```
CE#sh ip bgp
BGP table version is 3, local router ID is 192.168.2.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
x best-external, a additional-path, c RIB-compressed,
t secondary path,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

Network Next Hop Metric LocPrf Weight Path
* 1.1.1.1/32 192.168.2.2 0 111 111 200 100 i
*> 192.168.1.1 0 0 100 i
* 2.2.2.2/32 192.168.2.2 0 0 111 111 200 i
*> 192.168.1.1 0 100 200 i
CE#
```

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

The issue came after prepending AS 111 two times:

```
CE#sh ip bgp
BGP table version is 1, local router ID is 192.168.2.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
x best-external, a additional-path, c RIB-compressed,
t secondary path,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

Network Next Hop Metric LocPrf Weight Path
* 1.1.1.1/32 192.168.1.1 0 0 100 i
* 192.168.2.2 0 111 111 200 100 i
* 2.2.2.2/32 192.168.1.1 0 100 200 i
* 192.168.2.2 0 0 111 111 200 i
CE#
```

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Ignore this please, I just did not give enough time to bgp to complete the bgp table.

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

I labbed it, and do not see the issue with the disappearing routes. Original configuration. All routes are there in bgp table.

```
CE#sh ip bgp
BGP table version is 4, local router ID is 192.168.1.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
x best-external, a additional-path, c RIB-compressed,
t secondary path,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

Network Next Hop Metric LocPrf Weight Path
* 1.1.1.1/32 192.168.2.2 0 111 200 100 i
*> 192.168.1.1 0 0 100 i
* 2.2.2.2/32 192.168.1.1 0 100 200 i
*> 192.168.2.2 0 0 111 200 i
CE#
```

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

```
CE#sh run | sec router bgp
router bgp 64555
bgp log-neighbor-changes
neighbor 192.168.1.1 remote-as 100
neighbor 192.168.2.2 remote-as 200
neighbor 192.168.2.2 route-map AS in
CE#sh run | sec route-map
neighbor 192.168.2.2 route-map AS in
route-map AS permit 10
set as-path prepend 111
CE#
```

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Adding an as to the as-path only once did not change the path selection at all. They need to prepend two times: "set as-path prepend 111 111" or "set as-path prepend 200 200"

upvoted 1 times

drxz 7 months, 3 weeks ago

Selected Answer: B

When you want to prepend a route, you need the route-map to be "out" , not in.

upvoted 2 times

potato_inet0 7 months ago

not true, you can preepend AS-path as the route is coming in, so the inbound policy will preepend the AS then the route will be added to Local-RIB with the AS-path preepended, the question seems wrong in my opinion, none of the answers makes semse

upvoted 3 times

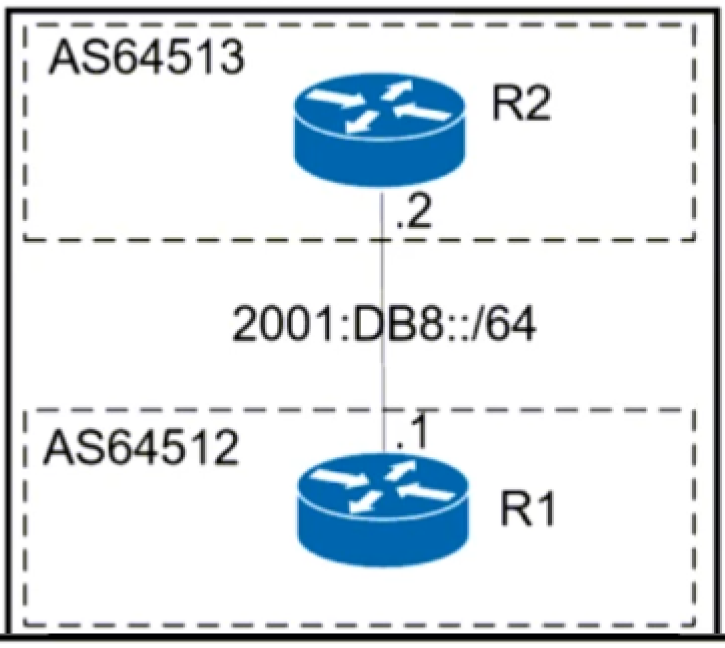

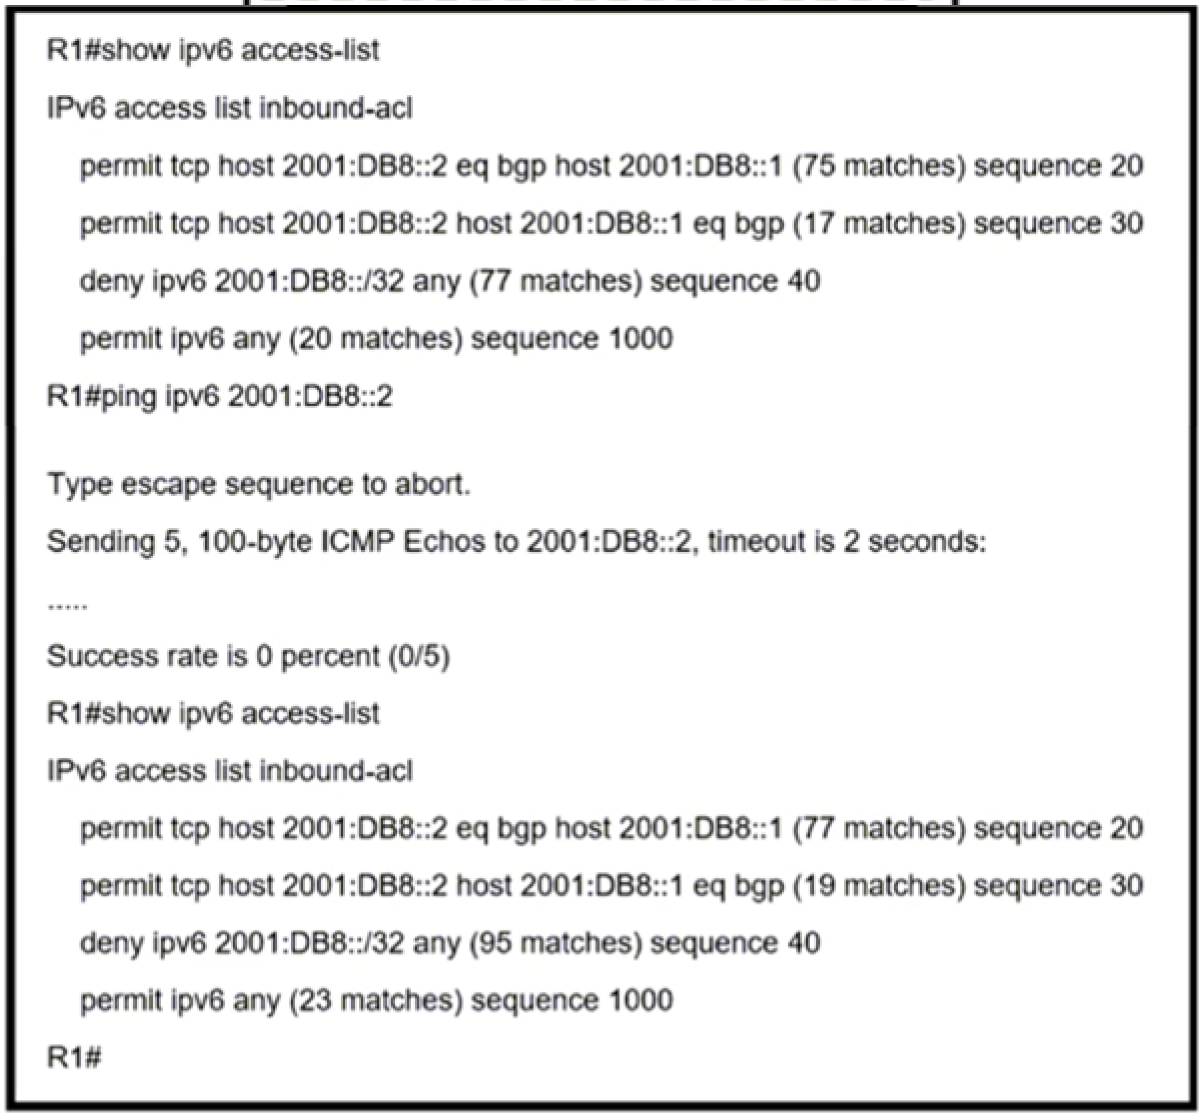

```
R1#show ipv6 access-list
IPv6 access list inbound-acl
   permit tcp host 2001:DB8::2 eq bgp host 2001:DB8::1 (75 matches) sequence 20
   permit tcp host 2001:DB8::2 host 2001:DB8::1 eq bgp (17 matches) sequence 30
   deny ipv6 2001:DB8::/32 any (77 matches) sequence 40
   permit ipv6 any (20 matches) sequence 1000
R1#ping ipv6 2001:DB8::2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8::2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R1#show ipv6 access-list
IPv6 access list inbound-acl
   permit tcp host 2001:DB8::2 eq bgp host 2001:DB8::1 (77 matches) sequence 20
   permit tcp host 2001:DB8::2 host 2001:DB8::1 eq bgp (19 matches) sequence 30
   deny ipv6 2001:DB8::/32 any (95 matches) sequence 40
   permit ipv6 any (23 matches) sequence 1000
R1#
```

Refer to the exhibit. An engineer applied an IPv6 traffic filter on R1. The interface flapped between R1 and R2 and clearing the BGP session did not restore the BGP session and failed. Which action must the engineer take to restore the BGP session from R2 to R1?

A. ICMPv6 must be permitted by the IPv6 traffic filter.

B. Swap the source and destination IP addresses in the IPv6 traffic filter.

C. Enable the BGP session, which went down when the session was cleared.

D. Apply the IPv6 traffic filter in the outbound direction on the interface.

Correct Answer: *A*

DeWalt95 2 weeks, 6 days ago

Honestly completely missed that BGP uses ICMP for keepalives.

upvoted 1 times

MasterMatt 4 months, 2 weeks ago

BGP use ICMP for keepalive. Thus why ICMPv6 is required.

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: A

Solution "A" makes sure that neighbor advertisement packets are not discarded. (Usually only required after adding "deny ipv6 any any").
https://alexandremspmoraes.wordpress.com/2012/02/17/how-ios-ipv6-acls-handle-icmpv6-neighbor-discovery-messages/

upvoted 3 times

inteldarvid 5 months, 1 week ago

Brothers u r really usefull, i hope if u have time to contact me here, i need to discuss some tasks with u kindly, +5491173651673 whats

upvoted 1 times

MEDO95 6 months, 2 weeks ago

Brother u r really usefull, i hope if u have time to contact me here, i need to discuss some tasks with u kindly, +15749999611 whats

upvoted 1 times

inteldarvid 5 months, 1 week ago

Brothers u r really usefull, i hope if u have time to contact me here, i need to discuss some tasks with u kindly, +5491173651673 whats

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

The connection never came up with source and destination swapped. So not "B".

```
R1(config)#do sh ip bgp all su
For address family: IPv6 Unicast
BGP router identifier 1.1.1.1, local AS number 1
BGP table version is 1, main routing table version 1

Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
2001:DB8::2 4 2 0 0 1 0 0 00:04:22 Active
R1(config)#do sh ipv6 access
IPv6 access list INBOUND
permit tcp host 2001:DB8::1 eq bgp host 2001:DB8::2 sequence 10
permit tcp host 2001:DB8::1 host 2001:DB8::2 eq bgp sequence 20
deny ipv6 2001:DB8::/32 any (131 matches) sequence 30
permit ipv6 any any (21 matches) sequence 40
R1(config)#
```

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

I applied this config in the lab (+reset bgp peers), and the result was that bgp dd not get blocked by this ACL. Am I missing something?

```
R1(config)#do sh ipv6 access-l
IPv6 access list INBOUND
permit tcp host 2001:DB8::2 eq bgp host 2001:DB8::1 (10 matches) sequence 10
permit tcp host 2001:DB8::2 host 2001:DB8::1 eq bgp sequence 20
deny ipv6 2001:DB8::/32 any sequence 30
permit ipv6 any any (1 match) sequence 40
R1(config)#
R1(config)#do sh run | sec GigabitEthernet0/0
interface GigabitEthernet0/0
ipv6 address 2001:DB8::1/64
ipv6 traffic-filter INBOUND in
R1(config)#
```

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

After applying the ACL bgp connection is still OK. Also received one prefix from neighbor.

```
R1(config)#do sh ip bgp all su
For address family: IPv6 Unicast
BGP router identifier 1.1.1.1, local AS number 1
BGP table version is 2, main routing table version 2
1 network entries using 168 bytes of memory
1 path entries using 108 bytes of memory
```

```
1/1 BGP path/bestpath attribute entries using 160 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 460 total bytes of memory
BGP activity 1/0 prefixes, 1/0 paths, scan interval 60 secs

Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
2001:DB8::2 4 2 9 9 2 0 0 00:04:39 1
```

upvoted 1 times

Question #463 Topic 1

DRAG DROP

-

Drag and drop the IPv6 First-Hop Security features from the left onto the definitions on the right.

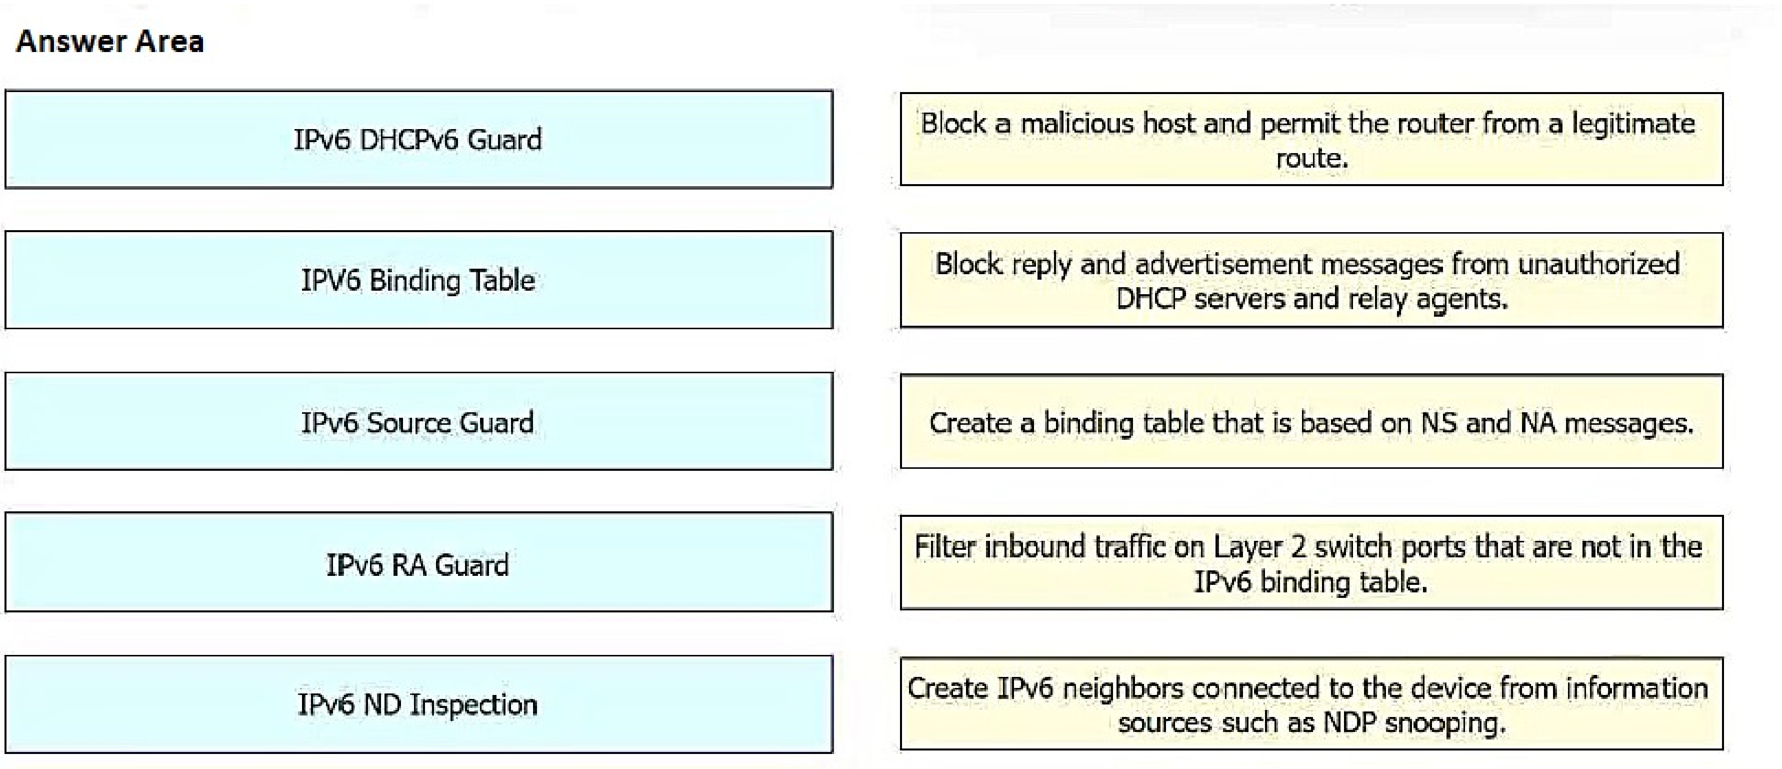

Answer Area

IPv6 DHCPv6 Guard	Block a malicious host and permit the router from a legitimate route.
IPV6 Binding Table	Block reply and advertisement messages from unauthorized DHCP servers and relay agents.
IPv6 Source Guard	Create a binding table that is based on NS and NA messages.
IPv6 RA Guard	Filter inbound traffic on Layer 2 switch ports that are not in the IPv6 binding table.
IPv6 ND Inspection	Create IPv6 neighbors connected to the device from information sources such as NDP snooping.

Correct Answer:

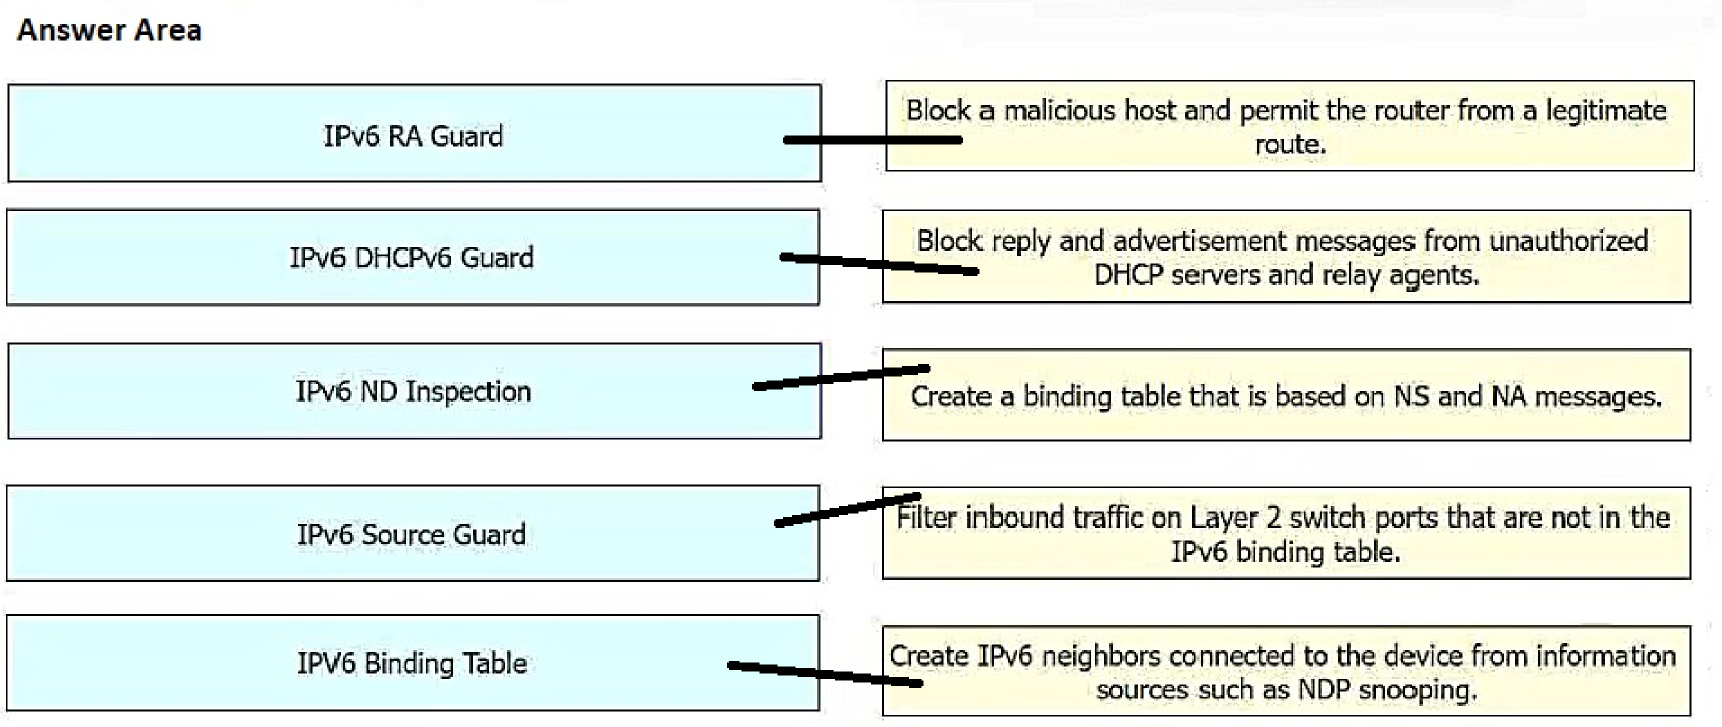

Answer Area

IPv6 RA Guard	Block a malicious host and permit the router from a legitimate route.
IPv6 DHCPv6 Guard	Block reply and advertisement messages from unauthorized DHCP servers and relay agents.
IPv6 ND Inspection	Create a binding table that is based on NS and NA messages.
IPv6 Source Guard	Filter inbound traffic on Layer 2 switch ports that are not in the IPv6 binding table.
IPV6 Binding Table	Create IPv6 neighbors connected to the device from information sources such as NDP snooping.

Fenix7 3 months, 1 week ago

Agree with seal2. Correct!

upvoted 2 times

seal2 3 months, 2 weeks ago

IPV6 DHCPV6 Guard -> Block reply and advertisement messages from unauthorized DHCP servers and relay agents.

IPV6 Binding Table -> Create IPV6 neighbors connected to the device from information
sources such as NDP snooping.

IPv6 Source Guard -> Filter inbound traffic on Layer 2 switch ports that are not in the IPV6 binding table.

IPv6 RA Guard -> Block a malicious host and permit the router from a legitimate route.

IPV6 ND Inspection -> Create a binding table that is based on NS and NA messages.

upvoted 3 times

inteldarvid 5 months, 1 week ago

correct

upvoted 1 times

Question #464

Topic 1

Refer to the exhibit.

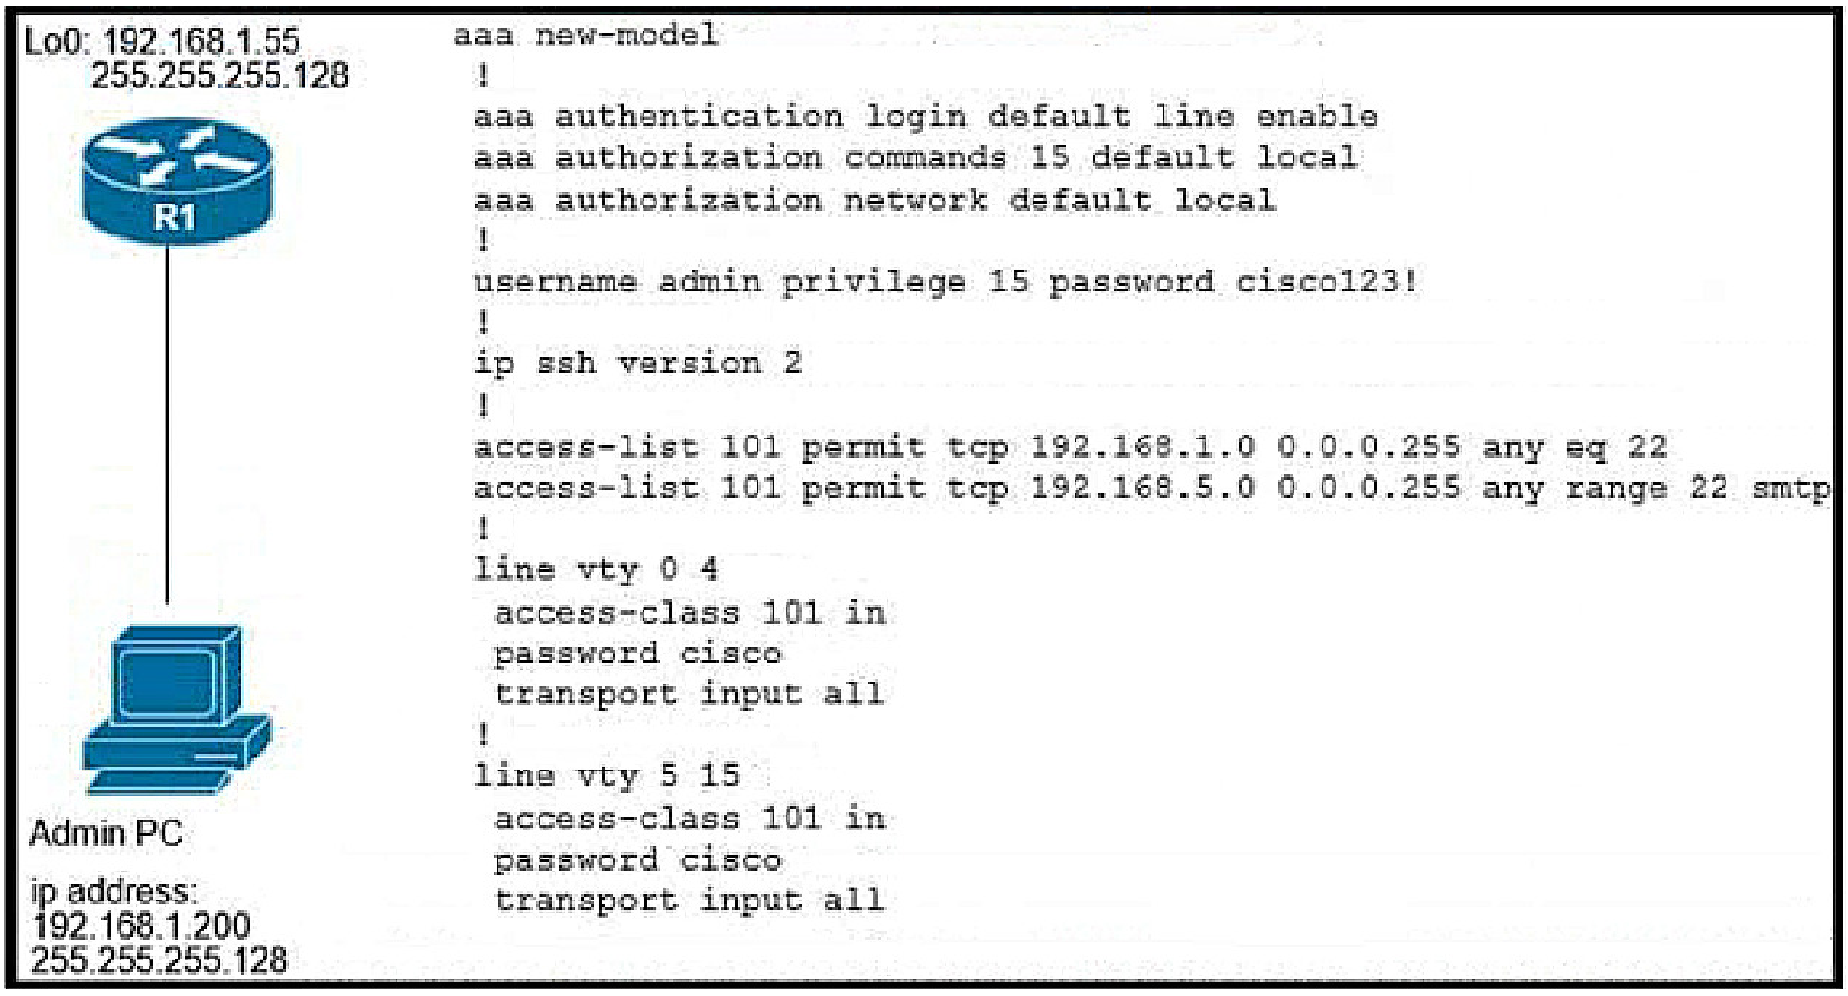

```
aaa new-model
 !
 aaa authentication login default line enable
 aaa authorization commands 15 default local
 aaa authorization network default local
 !
 username admin privilege 15 password cisco123!
 !
 ip ssh version 2
 !
 access-list 101 permit tcp 192.168.1.0 0.0.0.255 any eq 22
 access-list 101 permit tcp 192.168.5.0 0.0.0.255 any range 22 smtp
 !
 line vty 0 4
  access-class 101 in
  password cisco
  transport input all
 !
 line vty 5 15
  access-class 101 in
  password cisco
  transport input all
```

The administrator successfully logs into R1 but cannot access privileged mode commands. What should be configured to resolve the issue?

A. aaa authorization reverse-access

B. matching password on vty lines as cisco123!

C. secret cisco123! at the end of the username command instead of password cisco123!

D. enable secret or enable password commands to enter into privileged mode

Correct Answer: *D*

inteldarvid 5 months, 1 week ago

Selected Answer: D

correct

upvoted 1 times

potato_inet0 7 months, 2 weeks ago

D is the correct answer

upvoted 2 times

Question #465 Topic 1

Refer to the exhibit.

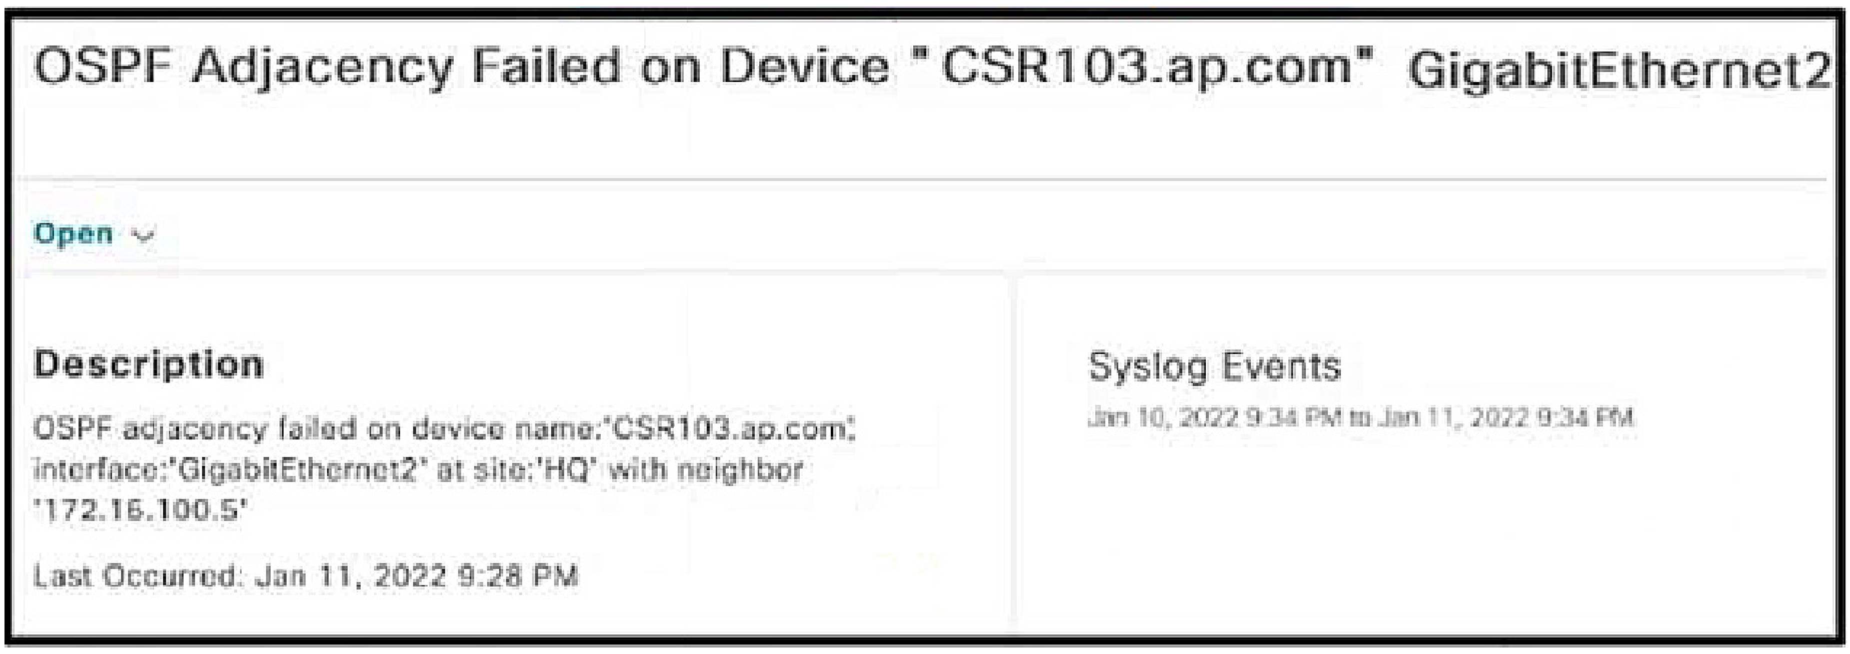

OSPF Adjacency Failed on Device " CSR103.ap.com" GigabitEthernet2

Open ⌄

Description

OSPF adjacency failed on device name:'CSR103.ap.com', interface:'GigabitEthernet2' at site:'HQ' with neighbor '172.16.100.5'

Last Occurred: Jan 11, 2022 9:28 PM

Syslog Events

Jan 10, 2022 9:34 PM to Jan 11, 2022 9:34 PM

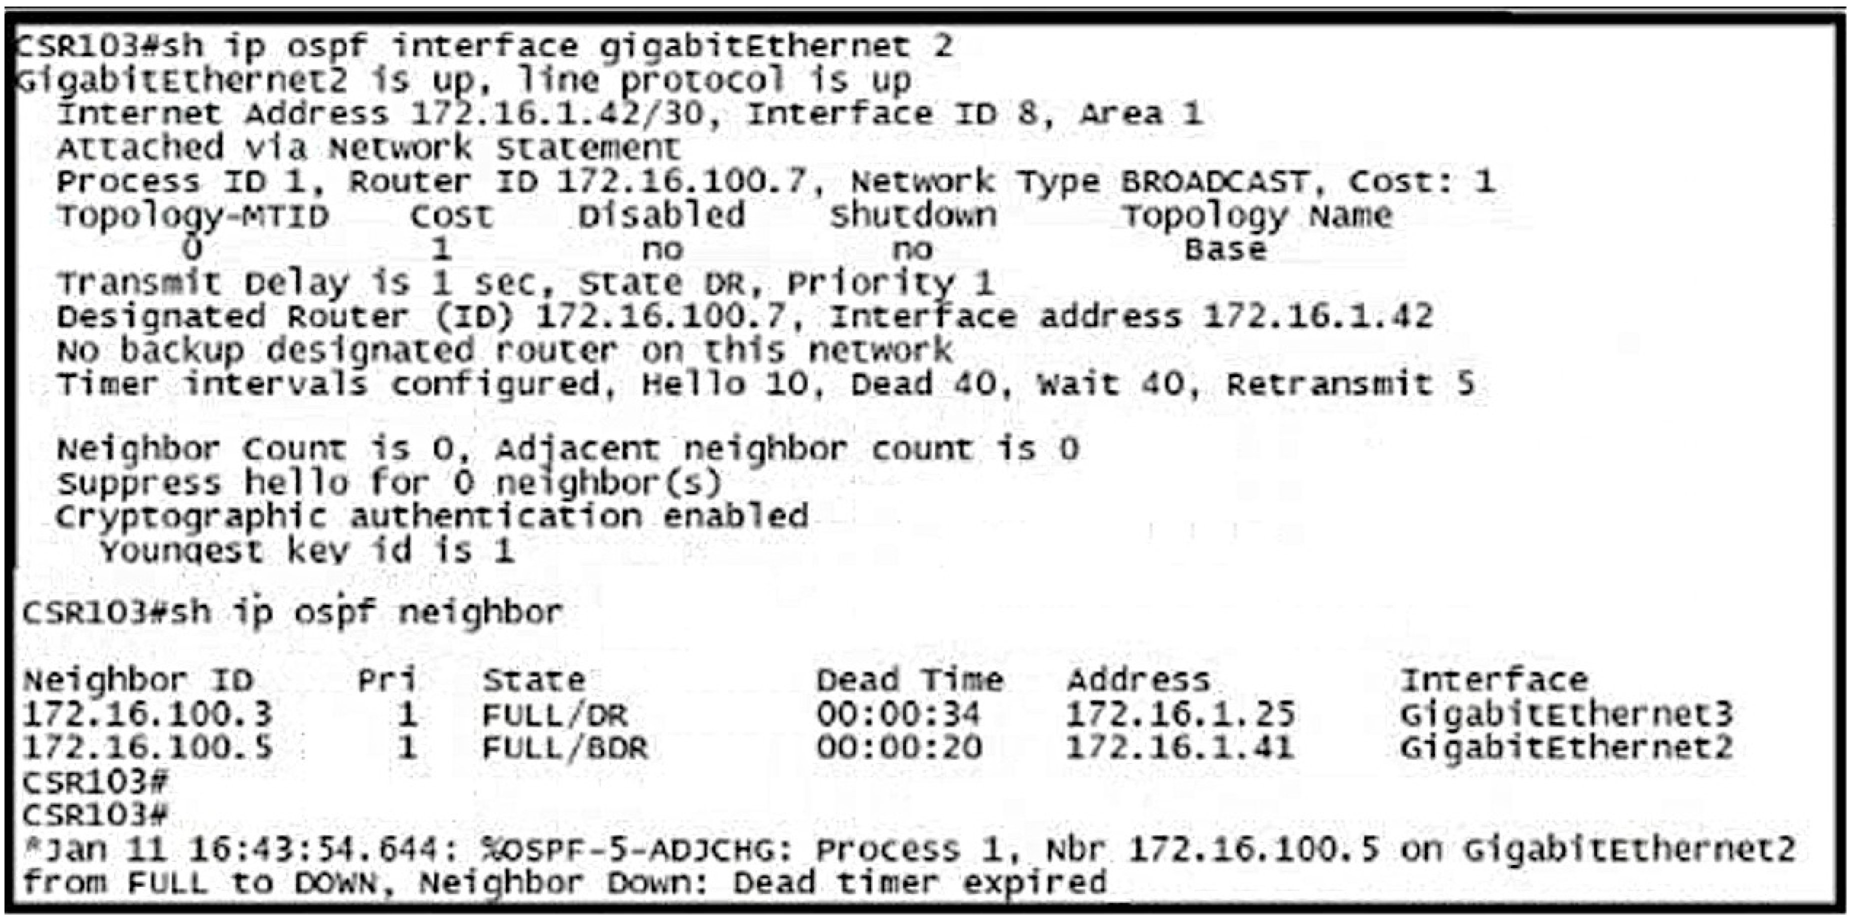

```
CSR103#sh ip ospf interface gigabitEthernet 2
GigabitEthernet2 is up, line protocol is up
  Internet Address 172.16.1.42/30, Interface ID 8, Area 1
  Attached via Network Statement
  Process ID 1, Router ID 172.16.100.7, Network Type BROADCAST, Cost: 1
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           1         no          no            Base
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 172.16.100.7, Interface address 172.16.1.42
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5

  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
  Cryptographic authentication enabled
    Youngest key id is 1

CSR103#sh ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.16.100.3      1   FULL/DR         00:00:34    172.16.1.25     GigabitEthernet3
172.16.100.5      1   FULL/BDR        00:00:20    172.16.1.41     GigabitEthernet2
CSR103#
CSR103#
*Jan 11 16:43:54.644: %OSPF-5-ADJCHG: Process 1, Nbr 172.16.100.5 on GigabitEthernet2
from FULL to DOWN, Neighbor Down: Dead timer expired
```

Which configuration must the engineer apply on CSR103 to resolve the problem?

A.

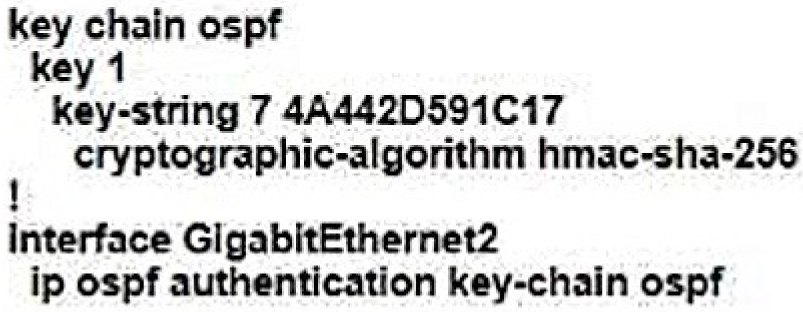

```
key chain ospf
 key 1
  key-string 7 4A442D591C17
   cryptographic-algorithm hmac-sha-256
!
interface GigabitEthernet2
 ip ospf authentication key-chain ospf
```

B.

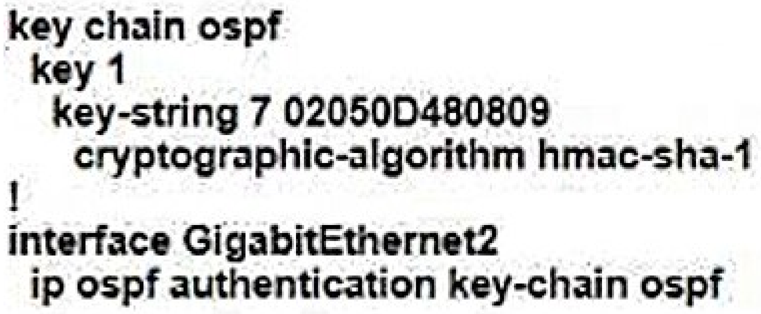

```
key chain ospf
 key 1
  key-string 7 02050D480809
   cryptographic-algorithm hmac-sha-1
!
interface GigabitEthernet2
 ip ospf authentication key-chain ospf
```

C. 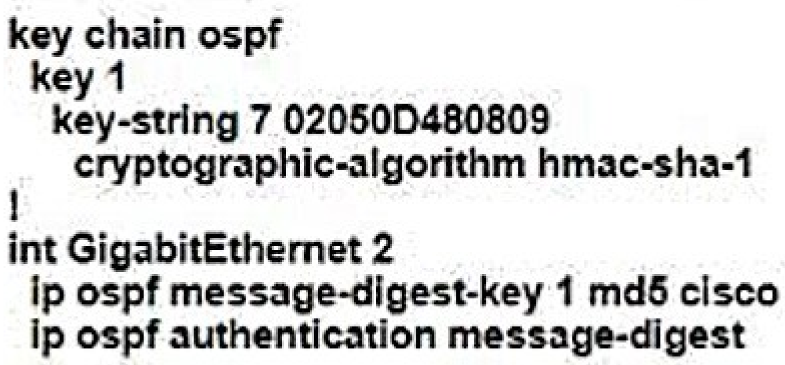

```
key chain ospf
 key 1
  key-string 7 02050D480809
   cryptographic-algorithm hmac-sha-1
!
int GigabitEthernet 2
 ip ospf message-digest-key 1 md5 cisco
 ip ospf authentication message-digest
```

D. 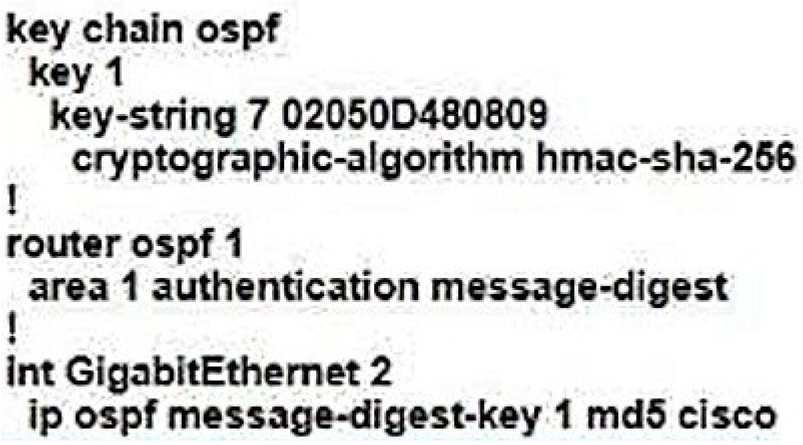

```
key chain ospf
 key 1
  key-string 7 02050D480809
   cryptographic-algorithm hmac-sha-256
!
router ospf 1
 area 1 authentication message-digest
!
int GigabitEthernet 2
 ip ospf message-digest-key 1 md5 cisco
```

Correct Answer: *B*

saiyuki1209 Highly Voted 8 months ago

Selected Answer: C

md5 authentication

Message digest authentication enabled
Youngest key id is 1
----
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13697-25.html

upvoted 6 times

HungarianDish Highly Voted 7 months ago

Selected Answer: C

The question is about ospf md5 authentication.
Based on the output md5 authentication has been enabled under the interface and not under the ospf process.
That's why it shows "Cryptographic authentication enabled" under "show ip ospf int gig 2".
This information is not displayed under the interface if authentication is enabled under the ospf process. (I labbed it.)
In this case, "C" is correct (and not "D", which enables it under the process.)

upvoted 5 times

HungarianDish 7 months ago

Example:
https://networklessons.com/ospf/how-to-configure-ospf-md5-authentication

interface GigabitEthernet 2
ip address 172.16.1.42 255.255.255.252
ip ospf 1 area 1
ip ospf authentication message-digest
ip ospf message-digest-key 1 md5 cisco

upvoted 2 times

HungarianDish 7 months ago

"A" and "B" are completly wrong.
The key chain configuration shown in the output is for eigrp, and the key chain was named as "ospf" to make the question tricky.

https://community.cisco.com/t5/switching/key-chain-md5-authentication-in-ospf/td-p/1327717
"OSPF is not using key chain, it is using authentication key you configured in the OSPF process or interface level."

upvoted 1 times

mouin Most Recent 3 months ago

Both C and D work, and the message "Youngest key id is 1" has nothing to do with the key chain.
I tried both (C&D) without configuring key chain and with key chain and they worked fine

upvoted 1 times

sal077 3 months, 4 weeks ago

Selected Answer: B

Not C or D because it's MD5, not Cryptograhpic as output states
Not A because the output should show SHA-256
So B it's correct.

It's IOS XE because is a CSR router:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/xe-16-10/iro-xe-16-10-book/iro-ospfv2-crypto-authen-xe.html

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: C

yes, option C

upvoted 1 times

Question #466 Topic 1

During the maintenance window, an administrator accidentally deleted the Telnet-related configuration that permits a Telnet connection from the inside network (Eth0/0) to the outside of the network between Friday - Sunday night hours only.

Which configuration resolves the issue?

A.

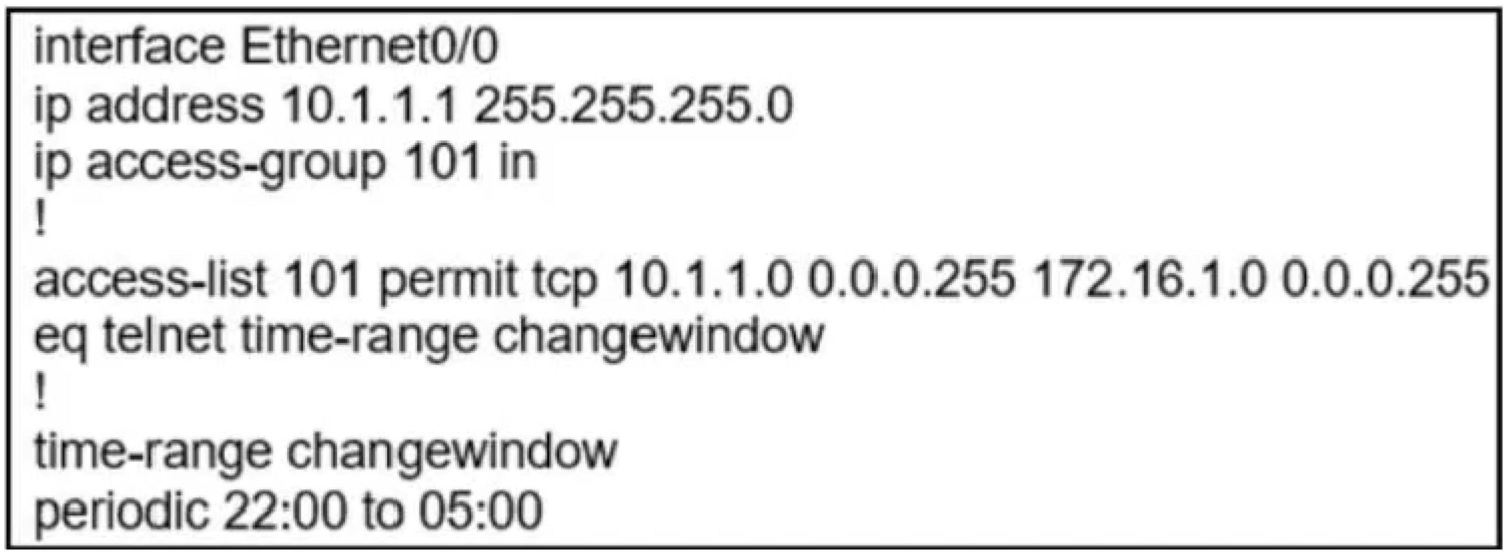

```
interface Ethernet0/0
ip address 10.1.1.1 255.255.255.0
ip access-group 101 in
!
access-list 101 permit tcp 10.1.1.0 0.0.0.255 172.16.1.0 0.0.0.255
eq telnet time-range changewindow
!
time-range changewindow
periodic 22:00 to 05:00
```

B.

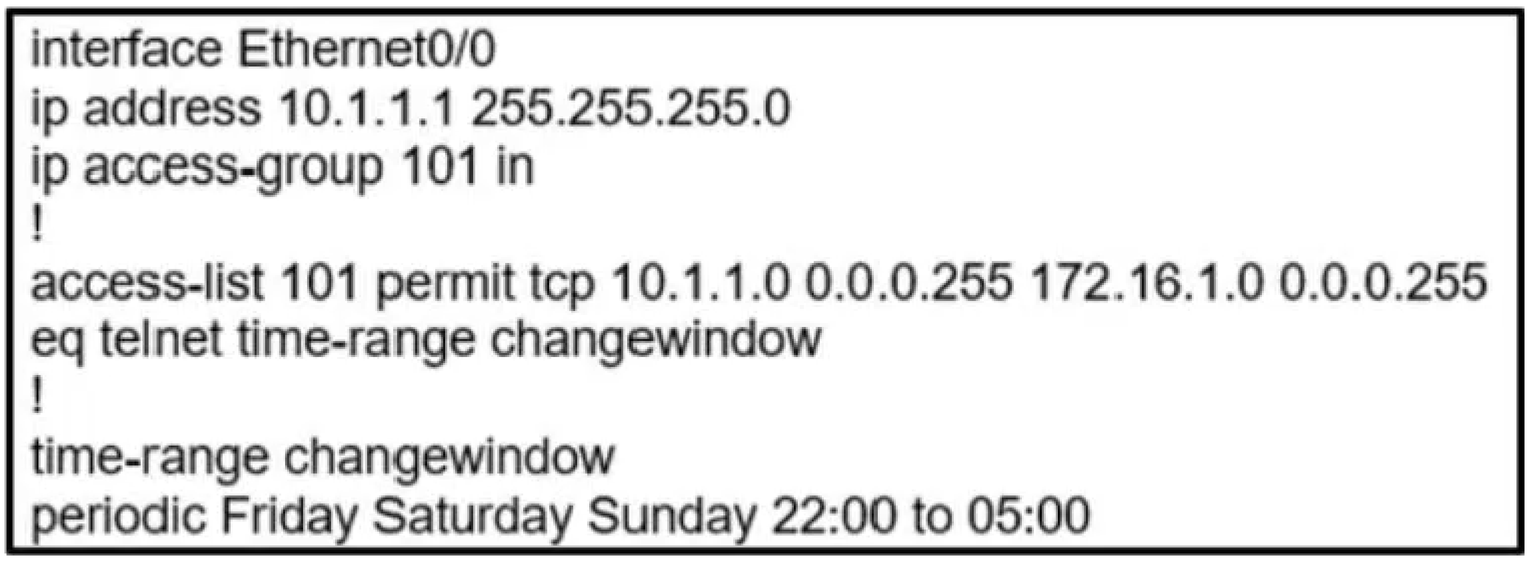

```
interface Ethernet0/0
ip address 10.1.1.1 255.255.255.0
ip access-group 101 in
!
access-list 101 permit tcp 10.1.1.0 0.0.0.255 172.16.1.0 0.0.0.255
eq telnet time-range changewindow
!
time-range changewindow
periodic Friday Saturday Sunday 22:00 to 05:00
```

C.

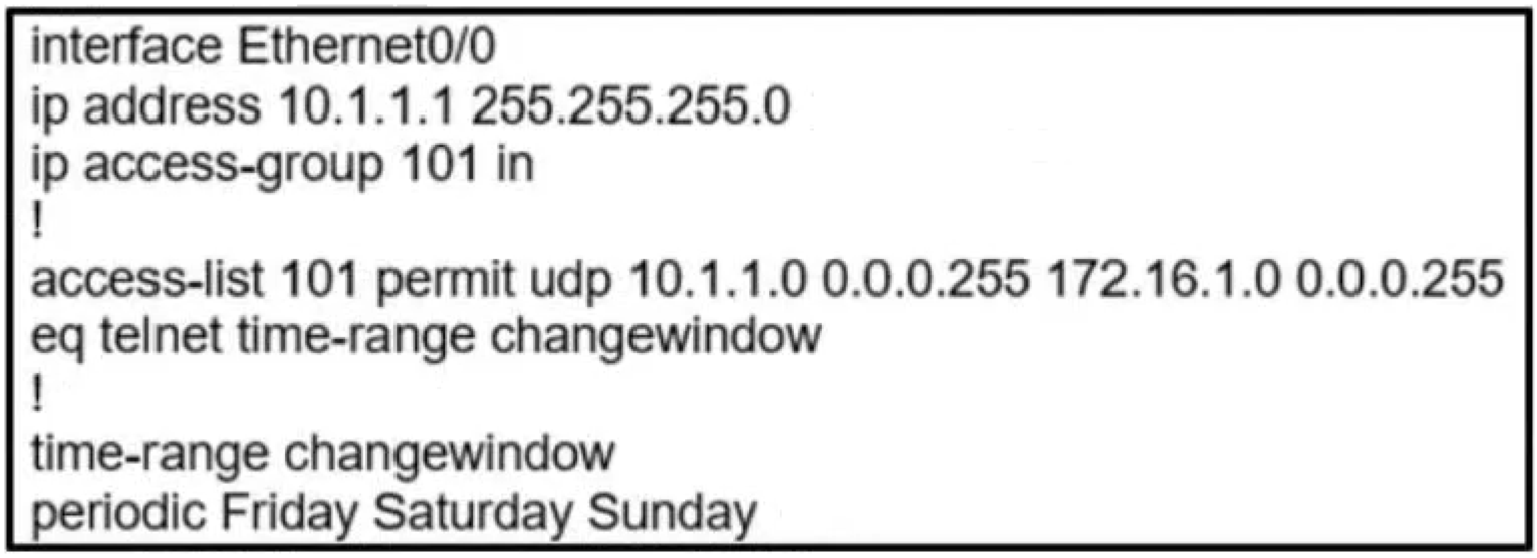

```
interface Ethernet0/0
ip address 10.1.1.1 255.255.255.0
ip access-group 101 in
!
access-list 101 permit udp 10.1.1.0 0.0.0.255 172.16.1.0 0.0.0.255
eq telnet time-range changewindow
!
time-range changewindow
periodic Friday Saturday Sunday
```

D.

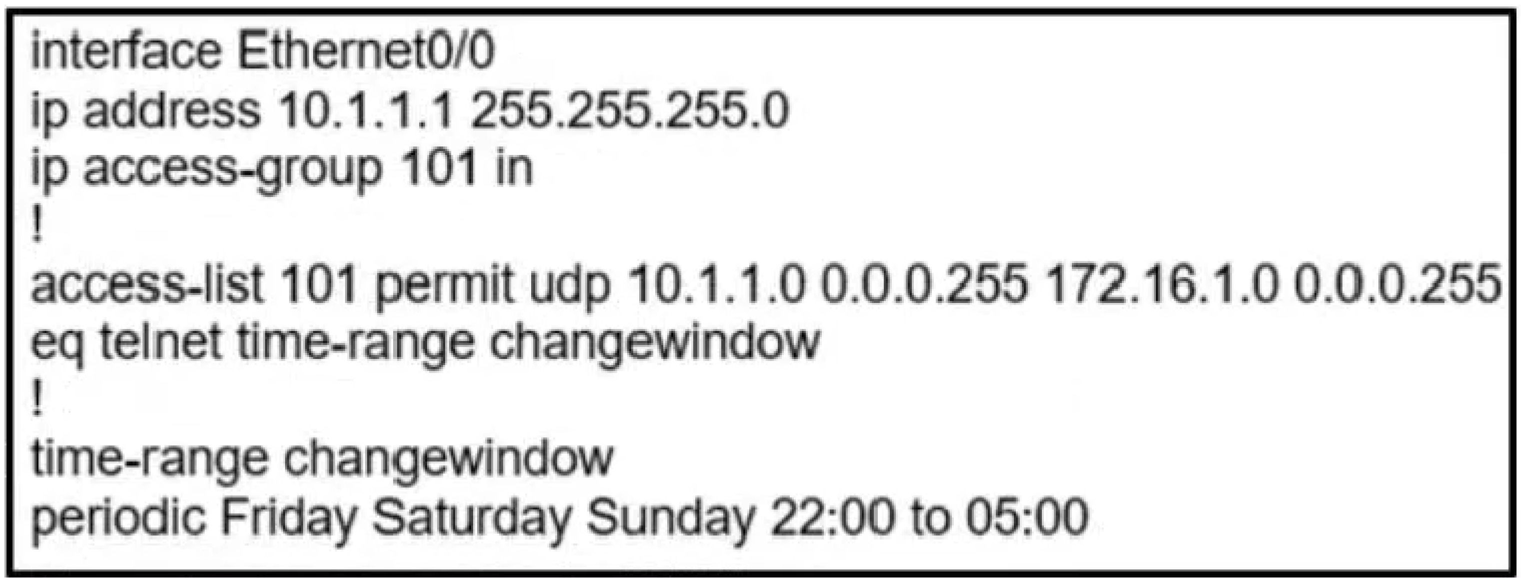

```
interface Ethernet0/0
ip address 10.1.1.1 255.255.255.0
ip access-group 101 in
!
access-list 101 permit udp 10.1.1.0 0.0.0.255 172.16.1.0 0.0.0.255
eq telnet time-range changewindow
!
time-range changewindow
periodic Friday Saturday Sunday 22:00 to 05:00
```

Correct Answer: *B*

⊟ **seal2** 3 months, 2 weeks ago

Selected Answer: B

B, telnet is TCP and it has friday thru sunday in the config

upvoted 2 times

robi1020 5 months ago

Selected Answer: B

Its B. TCP is the thing we are looking at.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: B

yes, b correct

upvoted 1 times

Question #467

Topic 1

SIMULATION

-

Guidelines

-

This is a lab item in which tasks will be performed on virtual devices.

- Refer to the Tasks tab to view the tasks for this lab item.
- Refer to the Topology tab to access the device console(s) and perform the tasks.
- Console access is available for all required devices by clicking the device icon or using the tab(s) above the console window.
- All necessary preconfigurations have been applied.
- Do not change the enable password or hostname for any device.
- Do not replace existing routing policies or configurations.
- Save your configurations to NVRAM before moving to the next item.
- Click Next at the bottom of the screen to submit this lab and move to the next question.
- When Next is clicked, the lab closes and cannot be reopened.

Topology

-

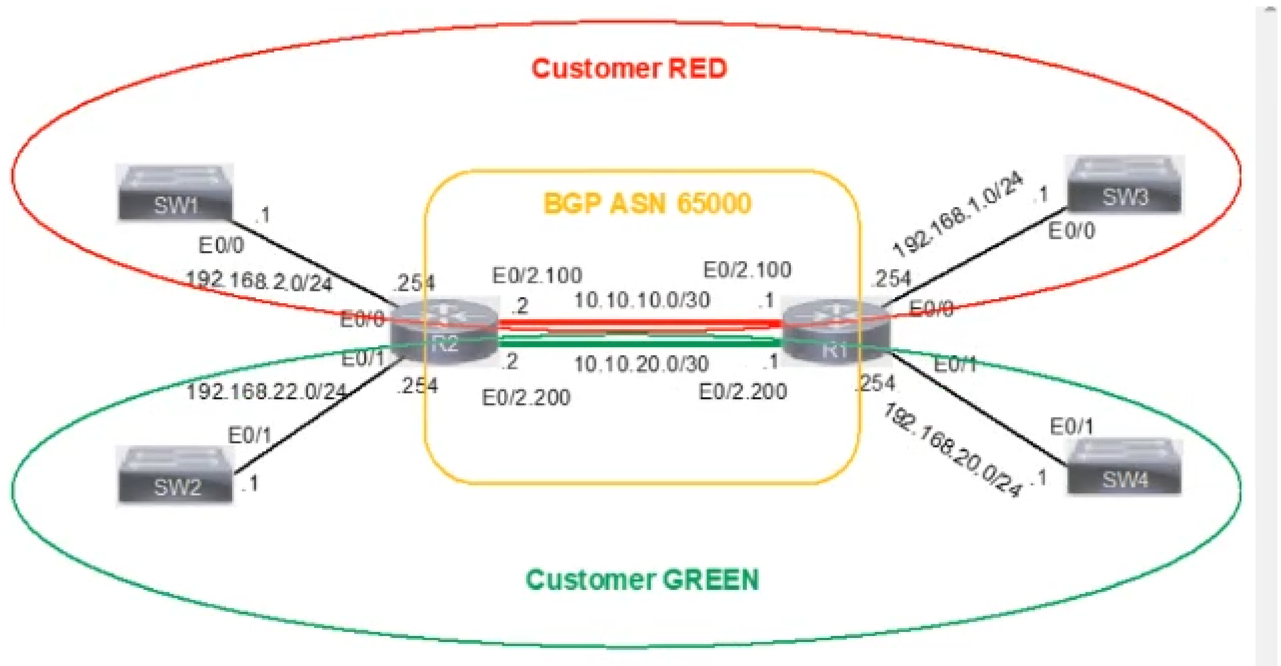

Topology Diagram

Tasks

-

Configure individual VRFs for each customer according to the topology to achieve these goals:

1. VRF “cu-red” has interfaces on routers R1 and R2. Both routers are preconfigured with IP addressing, VRFs, and BGP. Do not use the BGP network statement for advertisement.
2. VRF “cu-green” has interfaces on routers R1 and R2.
3. BGP on router R1 populates VRF routes between router R1 and R2.
4. BGP on router R2 populates VRF routes between router R1 and R2.
5. LAN to LAN is reachable between SW1 and SW3 for VRF “cu-red” and between SW2 and SW4 for VRF “cu-green”’. All switches are

preconfigured.

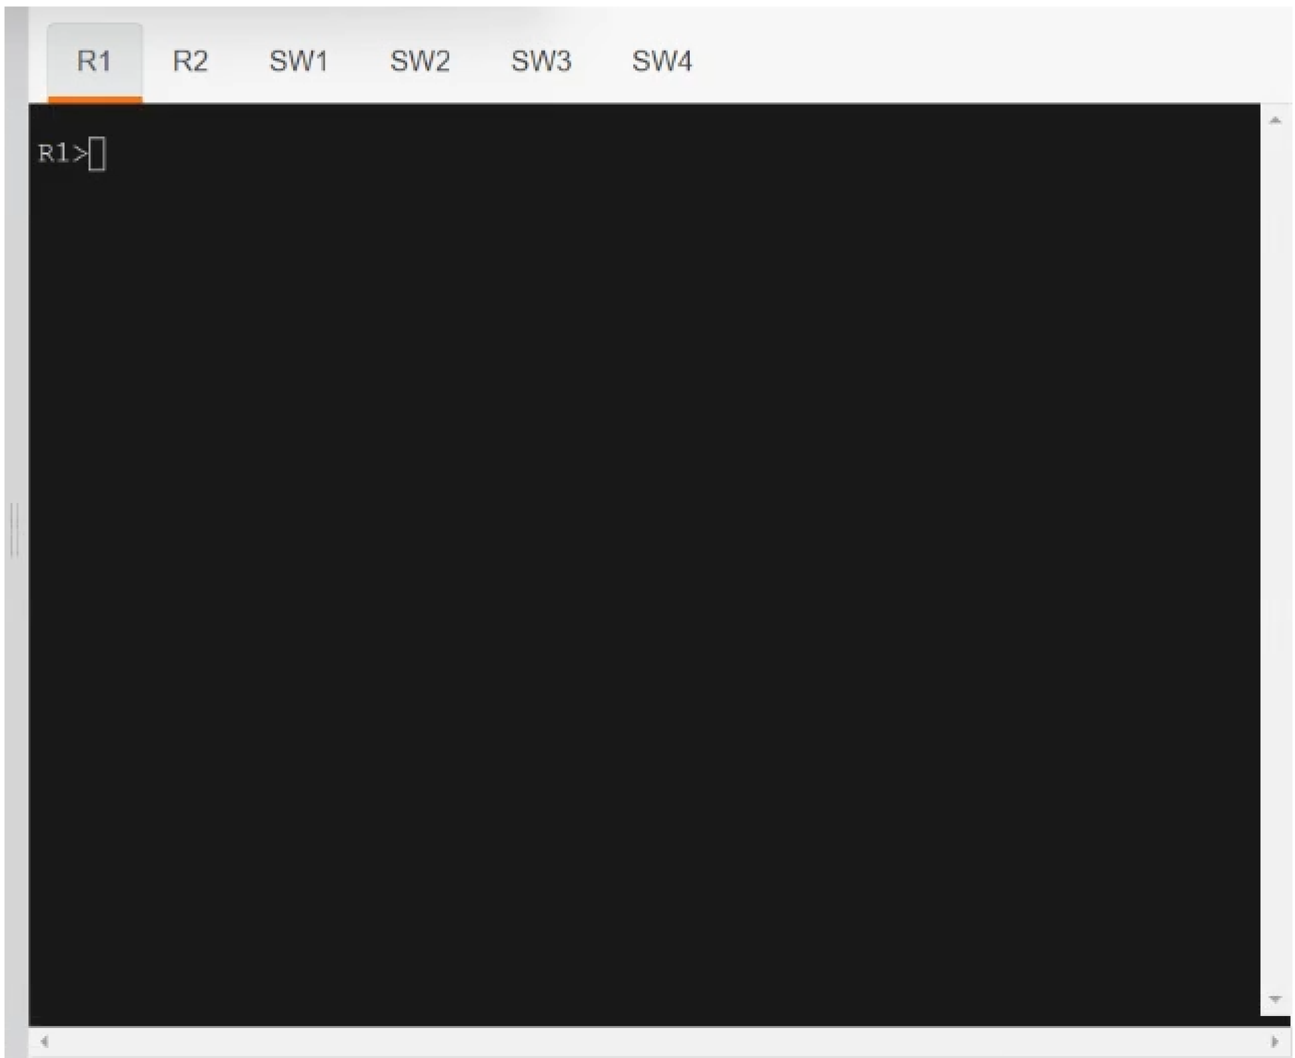

```
R1>
```

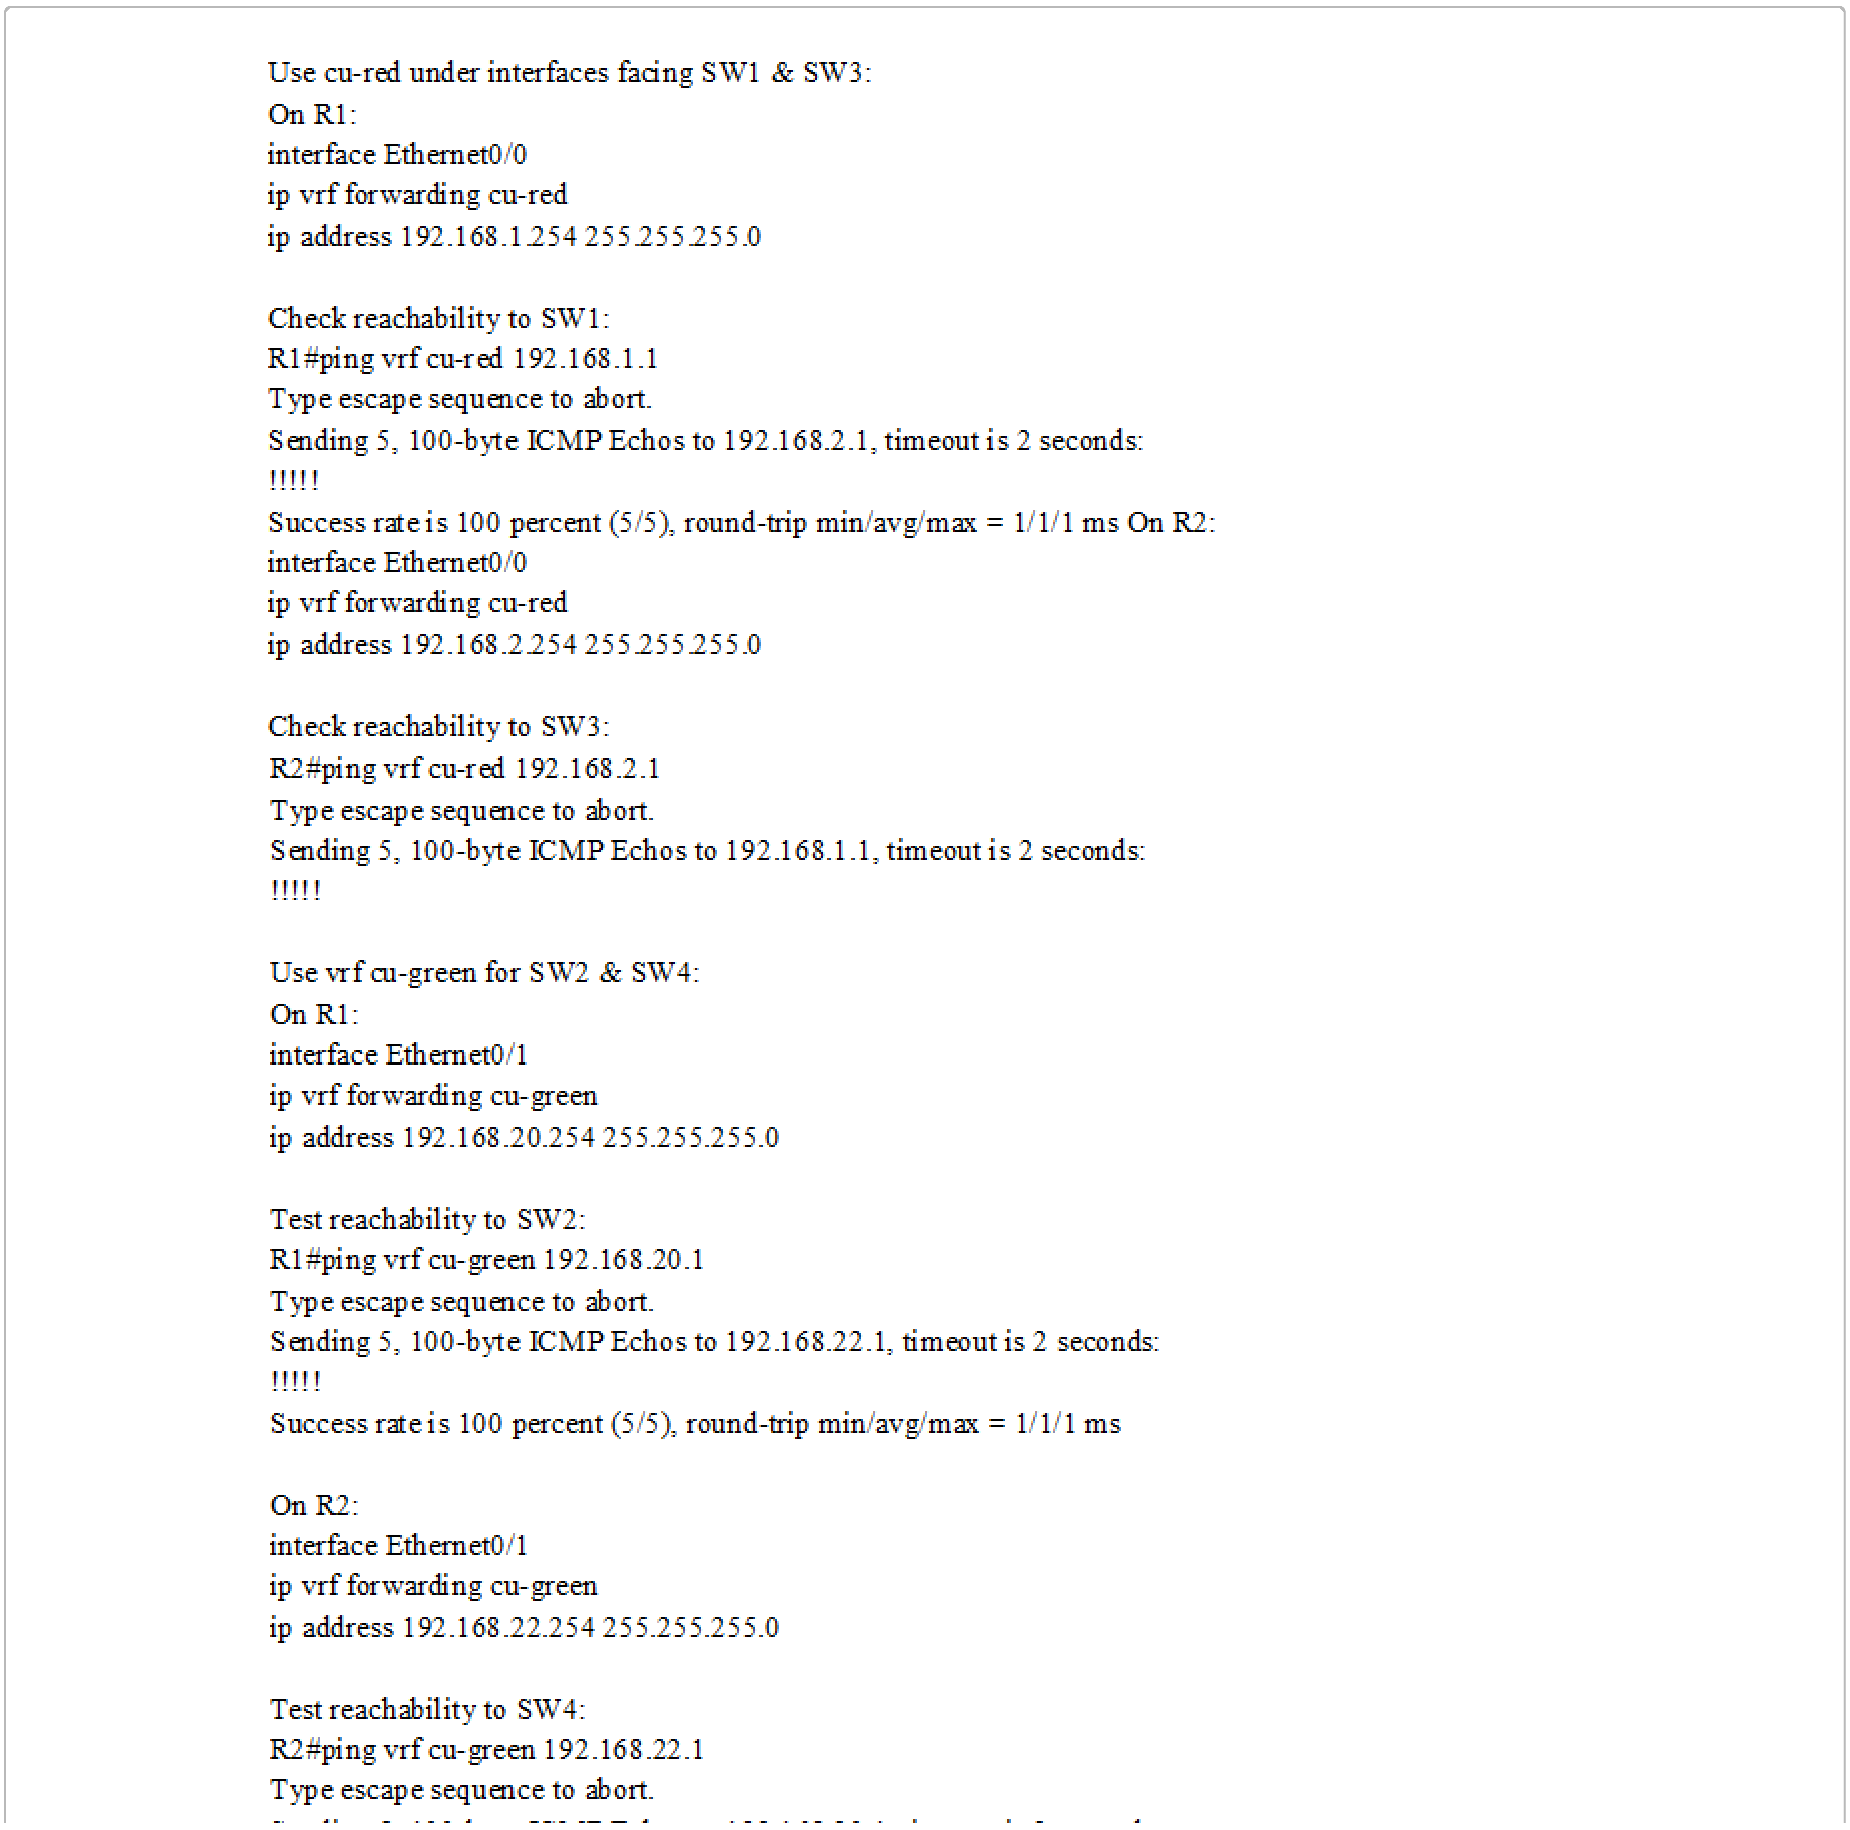

```
Use cu-red under interfaces facing SW1 & SW3:
On R1:
interface Ethernet0/0
ip vrf forwarding cu-red
ip address 192.168.1.254 255.255.255.0

Check reachability to SW1:
R1#ping vrf cu-red 192.168.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.2.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms On R2:
interface Ethernet0/0
ip vrf forwarding cu-red
ip address 192.168.2.254 255.255.255.0

Check reachability to SW3:
R2#ping vrf cu-red 192.168.2.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds:
!!!!!

Use vrf cu-green for SW2 & SW4:
On R1:
interface Ethernet0/1
ip vrf forwarding cu-green
ip address 192.168.20.254 255.255.255.0

Test reachability to SW2:
R1#ping vrf cu-green 192.168.20.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.22.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

On R2:
interface Ethernet0/1
ip vrf forwarding cu-green
ip address 192.168.22.254 255.255.255.0

Test reachability to SW4:
R2#ping vrf cu-green 192.168.22.1
Type escape sequence to abort.
```

Correct Answer:

```
Sending 5, 100-byte ICMP Echos to 192.168.20.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

On R1:
interface Ethernet0/2.100
mpls ip
!
interface Ethernet0/2.200
mpls ip
!
Configure BGP:
router bgp 65000
neighbor 10.10.10.2 remote-as 65000
neighbor 10.10.20.2 remote-as 65000
!
address-family vpnv4
neighbor 10.10.10.2 activate
neighbor 10.10.20.2 activate
exit-address-family
!
address-family ipv4 vrf cu-green
redistribute connected
exit-address-family
!
address-family ipv4 vrf cu-red
redistribute connected
exit-address-family
!
R1(config)#ip vrf cu-red
R1(config-vrf)#route-target both 65000:100
!
R1(config)#ip vrf cu-green
R1(config-vrf)#route-target both 65000:200

On R2:
interface Ethernet0/2.100
mpls ip
!
interface Ethernet0/2.200
mpls ip
!
router bgp 65000
neighbor 10.10.10.1 remote-as 65000
neighbor 10.10.20.1 remote-as 65000
!
address-family vpnv4
neighbor 10.10.10.1 activate
neighbor 10.10.20.1 activate
exit-address-family
!
address-family ipv4 vrf cu-green
redistribute connected
exit-address-family
!
address-family ipv4 vrf cu-red
redistribute connected
exit-address-family

R2(config)#ip vrf cu-red
R2(config-vrf)#route-target both 65000:100
!
R2(config)#ip vrf cu-green
R2(config-vrf)#route-target both 65000:200

Verification:
From SW1 to SW3:
SW1#ping 192.168.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
But can't Reach SW2 or SW4 in VRF cu-green:
SW1#ping 192.168.22.1
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.22.1, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)

SW1#ping 192.168.20.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.20.1, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)
```

Same Test for SW2:
From SW2 to SW4:

```
SW2#ping 192.168.20.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.20.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

But can’t Reach SW3 or SW1 in VRF cu-red:

```
SW2#ping 192.168.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)

SW2#ping 192.168.2.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.2.1, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)
```

Both R1 & R2 have separate tables for VRFs cu-red and cu-green.

DeWalt95 5 days, 6 hours ago

Labbed this and came to the same solution as Hungarian Dish.

upvoted 1 times

[Removed] 3 months, 2 weeks ago

Okay, so I did the following as the most simple solution I could think of.

```
Router-2
vrf definition cu-red
address-family ipv4 unicast
rd 65000:1
!
vrf definition cu-green
address-family ipv4 unicast
rd 65000:2
!
interface e0/0
description To SW-1
vrf forwarding cu-red
ip add 192.168.2.254 255.255.255.0
no shut
!
interface e0/1
description To SW-2
vrf forwarding cu-green
ip add 192.168.22.254 255.255.255.0
no shut
!
interface e0/2
description To R-1
no shut
!
interface e0/2.100
description To R-1 E0/2.100 CU-RED
vrf forwarding cu-red
ip address 10.10.10.2 255.255.255.252
no shut
!
interface e0/2.200
description To R-1 E0/2.200 CU-GREEN
vrf forwarding cu-green
```

```
ip address 10.10.20.2 255.255.255.252
no shut
!
router bgp 65000
bgp router-id 2.2.2.2
address-family ipv4 vrf cu-red
redistribute connected
neigh 10.10.10.10.1 remote-as 65000
address-family ipv4 vrf cu-green
redistribute connected
neigh 10.10.10.20.1 remote-as 65000
```

upvoted 3 times

[Removed] 3 months, 2 weeks ago

Continued...

```
Router-1
vrf definition cu-red
address-family ipv4 unicast
rd 65000:1
!
vrf definition cu-green
address-family ipv4 unicast
rd 65000:2
!
interface e0/0
description To SW-3
vrf forwarding cu-red
ip add 192.168.1.254 255.255.255.0
no shut
!
interface e0/1
description To SW-4
vrf forwarding cu-green
ip add 192.168.20.254 255.255.255.0
no shut
!
interface e0/2
description To R-2
no shut
!
interface e0/2.100
description To R-2 E0/2.100 CU-RED
vrf forwarding cu-red
ip address 10.10.10.1 255.255.255.252
no shut
!
interface e0/2.200
description To R-2 E0/2.200 CU-GREEN
vrf forwarding cu-green
ip address 10.10.20.1 255.255.255.252
no shut
!
router bgp 65000
bgp router-id 1.1.1.1
address-family ipv4 vrf cu-red
redistribute connected
neigh 10.10.10.10.2 remote-as 65000
address-family ipv4 vrf cu-green
redistribute connected
neigh 10.10.10.20.2 remote-as 65000
```

upvoted 2 times

[Removed] 3 months, 2 weeks ago

Switches should have a default-route to the R-1 or R-2 directly connected routed interface, for lab purposes. Not sure what the lab will be like.

```
SW-1
ip route 0.0.0.0 0.0.0.0 192.168.2.254
!
SW-2
ip route 0.0.0.0 0.0.0.0 192.168.22.254
!
SW-3
ip route 0.0.0.0 0.0.0.0 192.168.1.254
!
SW-4
ip route 0.0.0.0 0.0.0.0 192.168.20.254
```

upvoted 1 times

HungarianDish 6 months, 1 week ago

This sim was the same on the real exam, and got full score for it. All three labs (vrf, ospf, dmvpn) were evaluated.

upvoted 2 times

HungarianDish 6 months, 3 weeks ago

We do not need to set mpls ip, route target and bgp address-family vpnv4 for this. Just simple VRF-Lite (+ route distinguisher because the IOS prompts to use it, only locally significant).

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

```
R1
R1#sh run
ip vrf GREEN
rd 2:2
!
ip vrf RED
rd 1:1
!
interface GigabitEthernet0/0
ip vrf forwarding RED
ip address 192.168.1.254 255.255.255.0
!
interface GigabitEthernet0/1
ip vrf forwarding GREEN
ip address 192.168.20.254 255.255.255.0
!
interface GigabitEthernet0/2
no ip address
!
interface GigabitEthernet0/2.100
encapsulation dot1Q 100
ip vrf forwarding RED
ip address 10.10.10.1 255.255.255.252
!
interface GigabitEthernet0/2.200
encapsulation dot1Q 200
ip vrf forwarding GREEN
ip address 10.10.20.1 255.255.255.252
!
router bgp 65000
bgp router-id 1.1.1.1
bgp log-neighbor-changes
!
address-family ipv4 vrf GREEN
redistribute connected
neighbor 10.10.20.2 remote-as 65000
neighbor 10.10.20.2 activate
exit-address-family
!
address-family ipv4 vrf RED
redistribute connected
neighbor 10.10.10.2 remote-as 65000
neighbor 10.10.10.2 activate
exit-address-family
```

upvoted 5 times

HungarianDish 6 months, 3 weeks ago

Actually, VRF-lite is completely enough for this task. Example:
https://www.packetcoders.io/cisco-ios-how-to-configure-vrf-lite/

```
R1

R1#sh run
!
ip vrf green
rd 1:200
!
ip vrf red
rd 1:100
!
interface GigabitEthernet0/0
ip vrf forwarding red
ip address 192.168.1.254 255.255.255.0
!
interface GigabitEthernet0/1
ip vrf forwarding green
ip address 192.168.20.254 255.255.255.0
!
interface GigabitEthernet0/2
no ip address
!
interface GigabitEthernet0/2.100
encapsulation dot1Q 100
ip vrf forwarding red
ip address 10.10.10.1 255.255.255.252
!
interface GigabitEthernet0/2.200
encapsulation dot1Q 200
ip vrf forwarding green
ip address 10.10.20.1 255.255.255.252
!
router bgp 65000
bgp router-id 1.1.1.1
bgp log-neighbor-changes
!
address-family ipv4 vrf green
redistribute connected
neighbor 10.10.20.2 remote-as 65000
neighbor 10.10.20.2 activate
exit-address-family
!
address-family ipv4 vrf red
redistribute connected
neighbor 10.10.10.2 remote-as 65000
neighbor 10.10.10.2 activate
exit-address-family
```

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

```
R2#sh run
!
ip vrf green
rd 1:200
!
ip vrf red
rd 1:100
!
interface GigabitEthernet0/0
ip vrf forwarding red
ip address 192.168.2.254 255.255.255.0
!
interface GigabitEthernet0/1
ip vrf forwarding green
ip address 192.168.22.254 255.255.255.0
!
interface GigabitEthernet0/2
no ip address
!
interface GigabitEthernet0/2.100
encapsulation dot1Q 100
ip vrf forwarding red
ip address 10.10.10.2 255.255.255.252
!
interface GigabitEthernet0/2.200
encapsulation dot1Q 200
ip vrf forwarding green
ip address 10.10.20.2 255.255.255.252
!
router bgp 65000
bgp router-id 2.2.2.2
bgp log-neighbor-changes
!
address-family ipv4 vrf green
redistribute connected
neighbor 10.10.20.1 remote-as 65000
neighbor 10.10.20.1 activate
exit-address-family
!
address-family ipv4 vrf red
redistribute connected
neighbor 10.10.10.1 remote-as 65000
neighbor 10.10.10.1 activate
exit-address-family
```

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

```
R2#sh ip bgp all su
For address family: VPNv4 Unicast
BGP router identifier 2.2.2.2, local AS number 65000
BGP table version is 7, main routing table version 7
6 network entries using 936 bytes of memory
8 path entries using 672 bytes of memory
2/2 BGP path/bestpath attribute entries using 336 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1944 total bytes of memory
BGP activity 6/0 prefixes, 8/0 paths, scan interval 60 secs
```

```
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
10.10.10.1 4 65000 6 6 7 0 0 00:01:39 2
10.10.20.1 4 65000 5 5 7 0 0 00:00:38 2
```

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

```
sw1#ping 192.168.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/6 ms
sw1#ping 192.168.22.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.22.1, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)
sw1#ping 192.168.20.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.20.1, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)
sw1#
```

upvoted 1 times

HungarianDish 7 months ago

The provided solution has missing parts: route distinguisher for the vrf configuration, dot1q encapsulation for the subinterfaces, ip address for the mpls interfaces, extended community for carrying route-targets for MP-BGP and MPLS VPN. Plus, routing needs to be configured on the switches for successful testing.

upvoted 1 times

HungarianDish 7 months ago

I tested the scenario with this configuration in CML, and it worked:

```
R1
------------------------------------
ip vrf red
rd 1:100
route-target both 65000:100
ip vrf green
rd 1:200
route-target both 65000:200
int g0/0
ip vrf forwarding red
ip address 192.168.1.254 255.255.255.0
no shu
int g0/1
ip vrf forwarding green
ip address 192.168.20.254 255.255.255.0
no shu
int g0/2
no ip
no shut
int g0/2.100
enc dot1q 100
ip addr 10.10.10.1 255.255.255.252
mpls ip
int g0/2.200
enc dot1q 200
ip addr 10.10.20.1 255.255.255.252
mpls ip
router bgp 65000
bgp router 1.1.1.1
no synchronization
bgp log-neighbor-changes
no auto-summary
neigh 10.10.10.2 remote 65000
neigh 10.10.20.2 remote 65000
address-family vpnv4
neigh 10.10.10.2 activate
neigh 10.10.20.2 activate
neigh 10.10.10.2 send-community extended
neigh 10.10.20.2 send-community extended
address-family ipv4 vrf red
redist con
address-family ipv4 vrf green
redist con
```

upvoted 2 times

HungarianDish 7 months ago

```
R2
----------------------------------------
ip vrf red
```

```
rd 1:100
route-target both 65000:100
ip vrf green
rd 1:200
route-target both 65000:200
int g0/0
ip vrf forwarding red
ip address 192.168.2.254 255.255.255.0
no shu
int g0/1
ip vrf forwarding green
ip address 192.168.22.254 255.255.255.0
no shu
int g0/2
no ip
no shut
int g0/2.100
enc dot1q 100
ip addr 10.10.10.2 255.255.255.252
mpls ip
int g0/2.200
enc dot1q 200
ip addr 10.10.20.2 255.255.255.252
mpls ip
router bgp 65000
bgp router 2.2.2.2
no synchronization
bgp log-neighbor-changes
no auto-summary
neigh 10.10.10.1 remote 65000
neigh 10.10.20.1 remote 65000
address-family vpnv4
neigh 10.10.10.1 activate
neigh 10.10.20.1 activate
neigh 10.10.10.1 send-community extended
neigh 10.10.20.1 send-community extended
address-family ipv4 vrf red
redist con
address-family ipv4 vrf green
redist con
```

upvoted 2 times

⊟ **HungarianDish** 7 months ago

```
SW1 (red)
-----------
ip routing
int g0/0
no switchport
ip addr 192.168.2.1 255.255.255.0
no shu
vlan 100
ip route 0.0.0.0 0.0.0.0 192.168.2.254

SW2 (green)
------------
ip routing
int g0/1
no switchport
ip addr 192.168.22.1 255.255.255.0
no shu
vlan 200
ip route 0.0.0.0 0.0.0.0 192.168.22.254

SW3 (red)
------------
ip routing
int g0/0
no switchport
ip addr 192.168.1.1 255.255.255.0
no shu
vlan 100
ip route 0.0.0.0 0.0.0.0 192.168.1.254

SW4 (green)
-------------
ip routing
int g0/1
no switchport
ip addr 192.168.20.1 255.255.255.0
no shu
vlan 200
ip route 0.0.0.0 0.0.0.0 192.168.20.254
```

upvoted 3 times

Question #468 Topic 1

What is a function of BFD?

A. failure detection independent of routing protocols and media types

B. peer recovery after a Layer 2 adjacency failure

C. peer recovery after a Layer 3 protocol adjacency failure

D. failure detection dependent on routing protocols and media types

Correct Answer: *A*

seal2 3 months, 2 weeks ago

Selected Answer: A

correct, BFD detects failure without necessitating a specific routing protocol or media type.

upvoted 1 times

Refer to the exhibit.

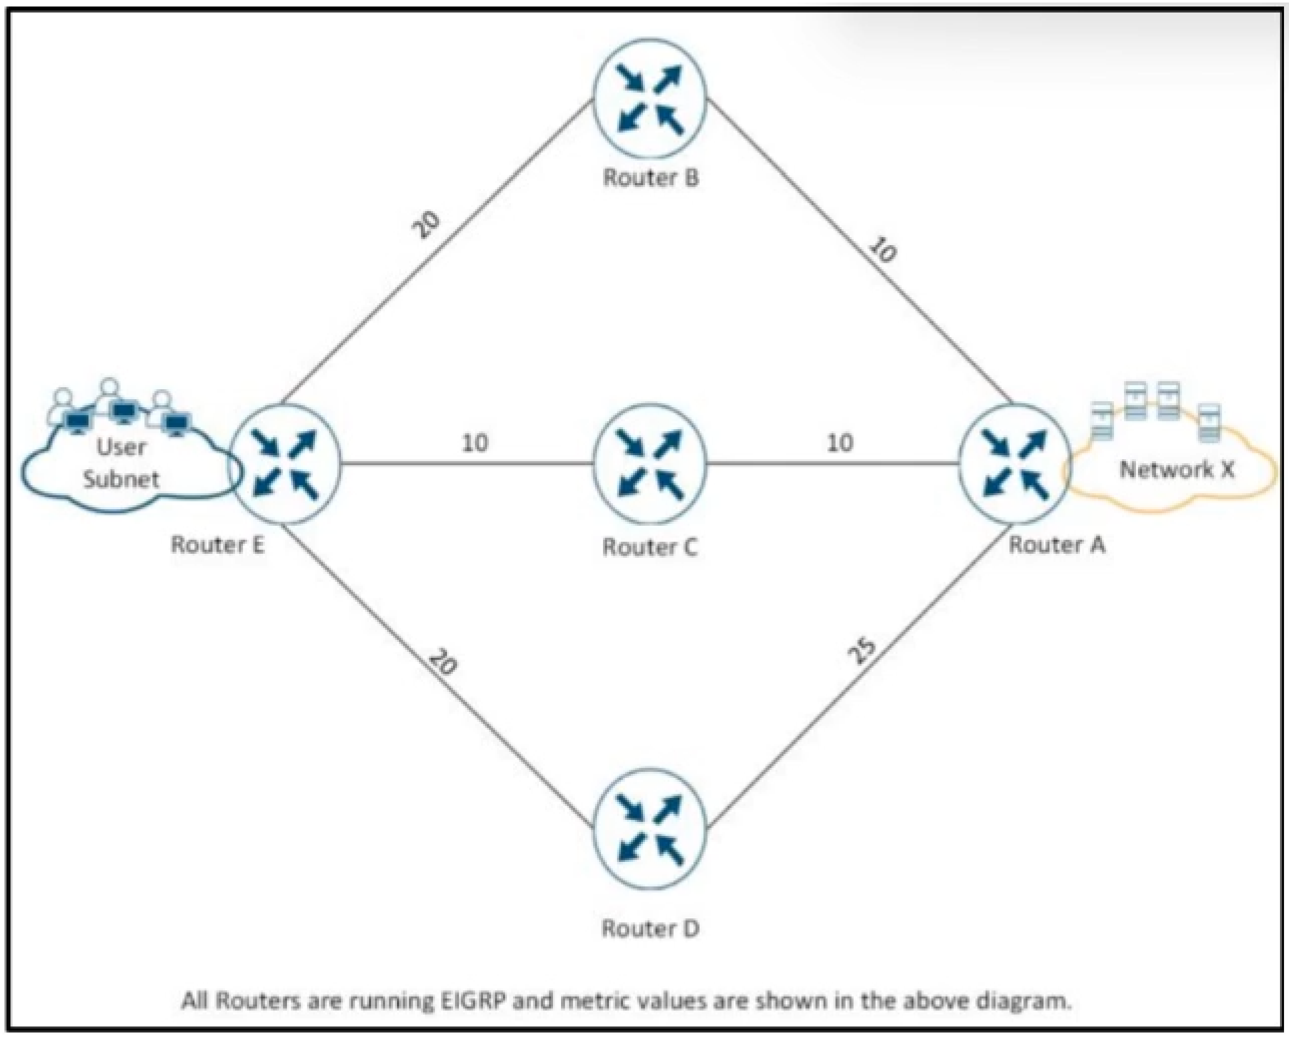
Router B
20
10
User Subnet
Router E
10
Router C
10
Router A
Network X
20
25
Router D
All Routers are running EIGRP and metric values are shown in the above diagram.

The IT manager received reports from users about slow applications through network x. Which action resolves the issue?

A. Upgrade the IOS on router E.

B. Move the servers into the users subnet.

C. Increase the bandwidth from the service provider.

D. Use the variance 2 command to enable load balancing.

Correct Answer: *D*

seal2 3 months, 2 weeks ago

Selected Answer: D

A. could potentially fix something... but very rarely. it's good to keep them updated but not because the older IOS's are necessarily slower.
B. what would moving them into the same subnet even do? you still have to actually send the data to the server through the network, needlessly complicating things more.
C. yeah lemme just ask my service provider for more bandwidth if things are slow. sounds reasonable. i mean, it would maybe be more plausible if we had reason to believe that was the problem but we don't, it appears that it's our own network causing issues.
D. load balancing is the most likely to actually do something and is something we can control. might as well try this first.

upvoted 1 times

HarwinderSekhon 4 months ago

is it correct? what about moving users to same subnet?

upvoted 1 times

Question #470 Topic 1

Refer to the exhibit.

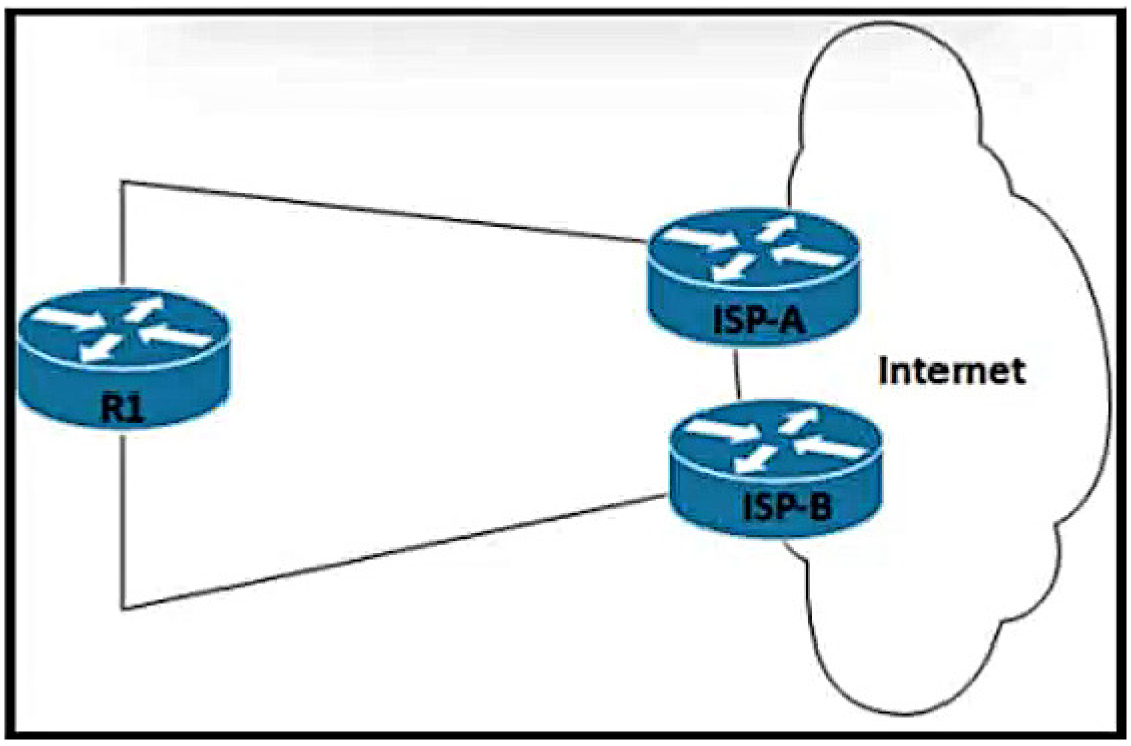

Router R1 peers with two ISPs using static routes to get to the internet. The requirement is that R1 must prefer ISP-A under normal circumstances and failover to ISP-B if the connectivity to ISP-A is lost. The engineer observes that R1 is load balancing traffic across the two ISPs.

Which action resolves the issue by sending traffic to ISP-A only with failover to ISP-B?

A. Configure two static routes on R1, one pointing to ISP-B with more specific routes and another pointing to ISP-A with summary routes.

B. Configure OSPF between R1, ISP-A, and ISP-B for dynamic failover if any ISP link to R1 fails.

C. Change the bandwidth of the interface on R1 so that interface to ISP-A has a higher value than the interface to ISP-B.

D. Configure two static routes on R1, one pointing to ISP-A and another pointing to ISP- B with 222 admin distance.

Correct Answer: *D*

sayed_2908 3 weeks ago

something is missing in the question. we should have the option to track if the link is down then route will automatically failover.

Adding two static route, even if the primary fails it won't get removed from the routing table will it?

upvoted 1 times

aqwsdfghjklp 1 month ago

I don't understand why B is incorrect.

upvoted 1 times

inteldarvid 5 months, 1 week ago

Selected Answer: D

Yes correct: option D, because route static float for example:
ISP1) ip route 0.0.0.0 0.0.0.0 200.0.0.1
isp2) ip route 0.0.0.0 0.0.0.0 90.0.0.1 222 (AD static float)

upvoted 1 times

cir_ 6 months, 3 weeks ago

D works if there is a layer 1 or 2 failure to ISP-A however I suspect there may be some details missing from this question

upvoted 1 times

GReddy2323 6 months, 4 weeks ago

I don't understand the reasoning behind answer D. Could someone provide more details?

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

https://networklessons.com/cisco/ccna-routing-switching-icnd1-100-105/floating-static-route
If we want to use a static route as a backup route, we'll have to change its administrative distance. This is called a floating static route.
https://www.geeksforgeeks.org/what-is-floating-static-route/
higher admin value to make a secondary route:

```
Router(config)#ip route 192.168.60.0 255.255.255.0 192.168.40.2 (primary route)
Router(config)#ip route 192.168.60.0 255.255.255.0 192.168.20.2 10 (secondary route)
```

upvoted 3 times

Question #471 *Topic 1*

SIMULATION

-

Guidelines

-

This is a lab item in which tasks will be performed on virtual devices.

- Refer to the Tasks tab to view the tasks for this lab item.
- Refer to the Topology tab to access the device console(s) and perform the tasks.
- Console access is available for all required devices by clicking the device icon or using the tab(s) above the console window.
- All necessary preconfigurations have been applied.
- Do not change the enable password or hostname for any device.
- Do not replace existing routing policies or configurations.
- Save your configurations to NVRAM before moving to the next item.
- Click Next at the bottom of the screen to submit this lab and move to the next question.
- When Next is clicked, the lab closes and cannot be reopened.

Topology

-

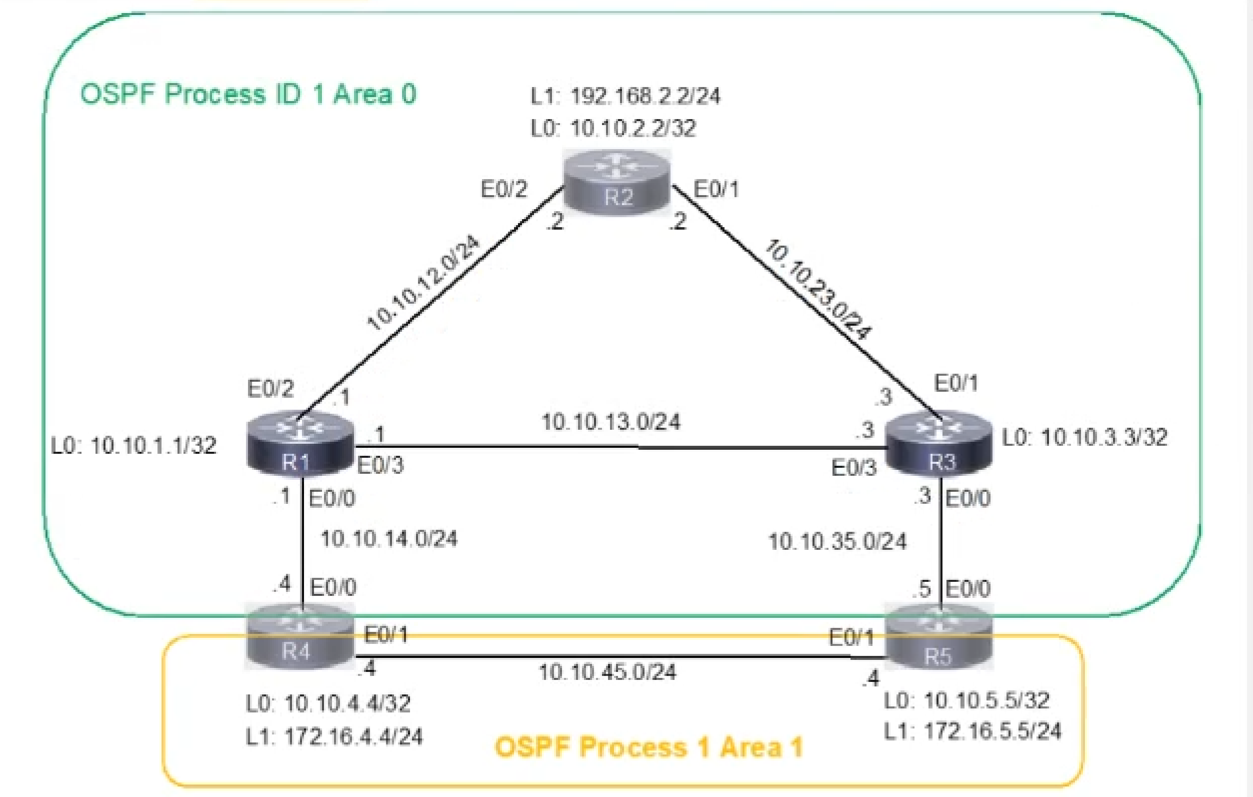

Topology Diagram

Tasks

-

A network is configured with IP connectivity, and the routing protocol between devices started having problems right after the maintenance window to implement network changes. Troubleshoot and resolve to a fully functional network to ensure that:

1. Inter-area links have link authentication (not area authentication) using MD5 with the key 1 string CCNP.
2. R3 is a DR regardless of R2 status while R1 and R2 establish a DR/BDR relationship.
3. OSPF uses the default cost on all interfaces. Network reachability must follow OSPF default behavior for traffic within an area over intra-area VS inter-area links.
4. The OSPF external route generated on R4 adds link cost when traversing through the network to reach R2. A network command to advertise

routes is not allowed.

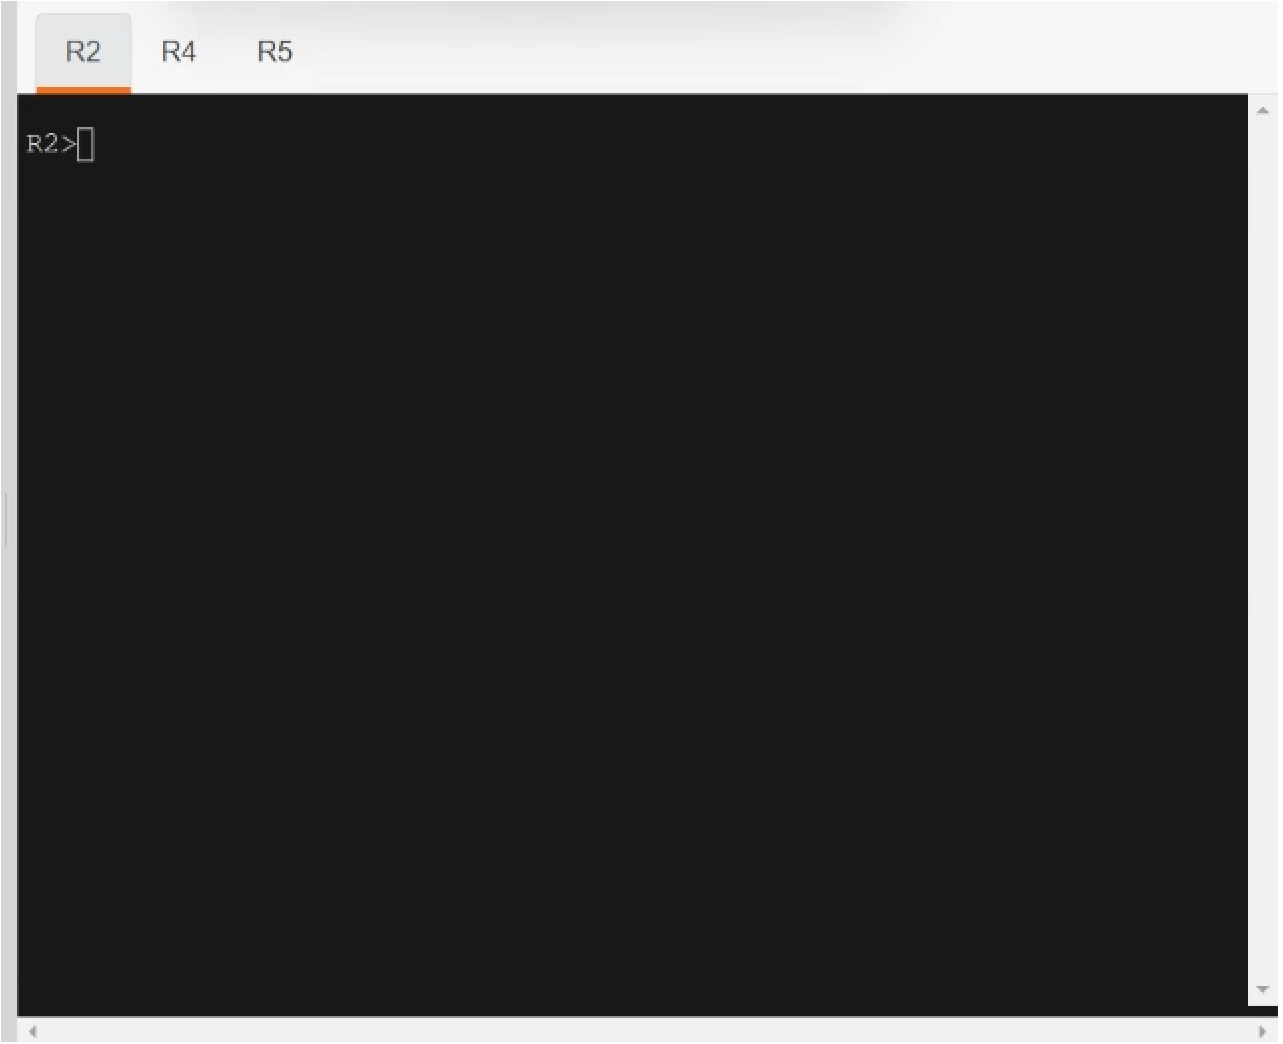

```
R2>
```

Correct Answer:

```
R4
int range et0/0 ?1
ip ospf authentication message-digest
ip ospf message-digest-key 1 md5 CCNP

Router ospf 1
Redistribute connected subnets route-map to-ospf metric-type 1

R5
int range et0/0 ?1
ip ospf authentication message-digest
ip ospf message-digest-key 1 md5 CCNP

interface eth 0/1
ip ospf cost 10
```

DeWalt95 4 days, 9 hours ago

Labbed out and came to the same solution as Hungarian Dish.

Do not need a route map to set R4s type 5 LSAs to metric-type 1 - as pointed out can be done in the redistribution statement.
Word of warning - setting the OSPF priority to 0 on any interface (as I tried initally) means the router will stay at DROTHER

upvoted 1 times

chris110 3 months, 1 week ago

OSPF LAB SOLUTION:

```
OSPF
R4# int e0/1
ip address 10.10.45.4
255.255.255.0
R4#
key chain CCNP key 1 key-string CCNP
Wr
R2# int e0/1
ip ospf priority 0
Wr
R5# int e0/1
no ip ospf cost 60
Wr
R4#
router ospf 1 redistribute connected subnets route-map to-ospf metric-type 1
```

upvoted 2 times

chris110 3 months, 1 week ago

BGP SIM:

A Single /16 is advertised for all infrastructure-connected interfaces that belong to the
10.20.x.x network using bgp network commands from border routers connected to the ISP (R4 & R5). DO NOT use the aggregte command.
R6 receives the ISP R2 loopback2 from R4 and receives a summary address for both loopbacks of ISP R2 from R4 and R5. Uses BGP-attribute local-preference add default value + router (ex R4 is local-preference 104) number. Use the existing prefix-list of route maps with the sequence numbering starting at 10 and added in increments of 10
R6 receives the ISP loopback1 from R5 and receives a summmary address for both loopbacks of
ISP R2 from R4 or R5 using the same guidelines.
R6 advertises it's loopback1 of 172.16.6.0/32 through BGP

upvoted 2 times

chris110 3 months, 1 week ago

Solution :

R4

```
R4(config)# ip route 10.20.0.0 255.255.0.0 null 0
R4(config)# no ip prefix-list AS65001-in
R4(config)# ip prefix-list AS65001-in seq 10 permit 192.168.2.0/24
R4(config-route-map)# route-map AS65001-in permit 20
R4(config-route-map)# set local-preference 104
R4(config-route-map)# end
R4# clear ip bgp * soft
```

upvoted 1 times

chris110 3 months, 1 week ago

Solution:

R5

```
R5(config)# ip route 10.20.0.0 255.255.0.0 null 0
R5(config)# no ip prefix-list AS65001-in
R5(config)# ip prefix-list AS65001-in seq 10 permit 192.168.3.0/24
R5(config-route-map)# route-map AS65001-in permit 20
R5(config-route-map)# set local-preference 105
R5(config-route-map)# end
R5# clear ip bgp * soft
```

upvoted 1 times

chris110 3 months, 1 week ago

Solution:

R6

```
router bgp 65000
address-family ipv4
no network 172.16.6.0
network 172.16.6.0 255.255.255.0
```

upvoted 2 times

Brand 3 months, 2 weeks ago

Guys I had the BGP version of this sim. There is a R6 connected to R4 and R5 as part of customer network. And they are running iBGP between R4-R5-R6.
R1-R2-R3 also running iBGP. R2 has 3 loopbacks. R6 should learn R2 Loopback2 IP from R4 and R2 Loopback1 from R5. Also, R6 should receive a summary route for both R2's L1-L2 from either R4 or R5 (you simply decide that). You are only allowed to configure R4/R5/R6. There is a prefix-list in both R4 and R5 for R2 loopbacks. And there is a route-map in R4 and R5 using the prefix-list to control prefix advertisements toward R6.

I can't recall all the details but I hope this helps. I think if you lab the above scenario and familiarize yourself with route-map/prefix-list to control BGP advertisements, then you'll be just fine.

upvoted 2 times

Brand 3 months, 2 weeks ago

By the way you're not allowed to use "aggregate" command in BGP so the route summarization for R2 L1-L2 should be done using prefix-list.

upvoted 2 times

alex711 3 months, 2 weeks ago

Thanks for sharing Brand!

Can you tell, which other sim you had ?

upvoted 1 times

Brand 3 months, 1 week ago

I had CoPP but I had to skip it as there was not enough time to complete it. The other was DMVPN. Check my comments on the DMVPN question.

upvoted 2 times

alex711 3 months, 1 week ago

Ok. Thanks again.
upvoted 1 times

⊟ **Chiaretta** 5 months, 1 week ago

Some points of this lab are very vague.
upvoted 1 times

⊟ **bizzar777** 5 months, 4 weeks ago

R2
int e0/1
ip ospf priority 0
R3 is a DR regardless of R2 status -> so is usefull to set ip ospf priority 0 on R2 E0/1
upvoted 3 times

⊟ **HungarianDish** 6 months, 1 week ago

This sim was bit different on the real exam, but got full score for it. All three labs (vrf, ospf, dmvpn) were evaluated. Only R2, R4, R5 are accessible. Need to add priority 0 on R2 for DR election (under interface towards R3). No route-map was needed.
upvoted 1 times

⊟ **ParisaAlipoor** 5 months, 2 weeks ago

did you pass the exam? for simulation exactly same? how many simulation question did you have?
upvoted 1 times

⊟ **HungarianDish** 5 months ago

Passed. 3 sims, all sims evaluated. Exact same sims as here.
upvoted 1 times

⊟ **spada05** 5 months ago

I tested today and had 2 different sims from what is here. The only matching was DMVPN. I had one for BGP and one for CoPP... Sure wish they were available here.
upvoted 2 times

⊟ **Stylar** 4 months, 3 weeks ago

Hello, do you recall the BGP sim ? What was there to configure? RR setup or some basics?
upvoted 1 times

⊟ **HungarianDish** 7 months ago

The first task is vague: "Inter-area links have link authentication". What do they mean by inter-area links? I only know about inter-area routes.
https://community.cisco.com/t5/switching/ospf-difference-inter-area-and-intra-area-routes/td-p/1900023
"Inter-area routes are the "O IA" routes that are learned in different areas."
upvoted 1 times

⊟ **HungarianDish** 7 months ago

I tested this scenario in CML with a somewhat different configuration. I think that we need to ensure that R3 is the DR on int g0/1 with a config like this: "ip ospf priority 255". Plus, we can use a simpler configuration on R4 for setting the metric type to E1:
router ospf 1
redistribute connected subnets metric-type 1
upvoted 1 times

⊟ **HungarianDish** 7 months ago

Well, this one works too, but I like to keep things simple:
router ospf 1
redistribute connected subnets route-map E1

route-map E1 permit 10
set metric-type type-1
upvoted 2 times

⊟ **HungarianDish** 7 months ago

```
R1
int lo0
ip addr 10.10.1.1 255.255.255.255
ip ospf 1 area 0

int g0/2
ip addr 10.10.12.1 255.255.255.0
no shu
ip ospf 1 area 0

int g0/3
ip addr 10.10.13.1 255.255.255.0
no shu
ip ospf 1 area 0

int g0/0
ip addr 10.10.14.1 255.255.255.0
no shu
ip ospf 1 area 0
ip ospf message 1 md5 CCNP
ip ospf auth message

router ospf 1
exi

R2
int lo0
ip addr 10.10.2.2 255.255.255.255
ip ospf 1 area 0
int lo 1
ip addr 192.168.2.2 255.255.255.0
ip ospf 1 area 0

int g0/2
ip addr 10.10.12.2 255.255.255.0
ip ospf 1 area 0
no shu
int g0/1
ip addr 10.10.23.2 255.255.255.0
ip ospf 1 area 0
no shu

router ospf 1
exi

R3

int lo 0
ip addr 10.10.3.3 255.255.255.255
ip ospf 1 area 0

int g0/1
ip addr 10.10.23.3 255.255.255.0
ip ospf 1 area 0
ip ospf priority 255
no shu

int g0/3
ip addr 10.10.13.3 255.255.255.0
ip ospf 1 area 0
no shu

int g0/0
ip addr 10.10.35.3 255.255.255.0
ip ospf 1 area 0
ip ospf message 1 md5 CCNP
ip ospf auth message
no shu

router ospf 1
exi
```

upvoted 1 times

⊟ **HungarianDish** 7 months ago

```
R4
int lo 0
ip addr 10.10.4.4 255.255.255.255
int lo 1
ip addr 172.16.4.4 255.255.255.0

int g0/0
ip addr 10.10.14.1 255.255.255.0
no shu
ip ospf 1 area 0
ip ospf message 1 md5 CCNP
ip ospf auth message

int g0/1
ip addr 10.10.45.4 255.255.255.0
no shu
ip ospf 1 area 1

router ospf 1
redistribute connected subnets metric-type 1

R5
int lo 0
ip addr 10.10.5.5 255.255.255.255
int lo 1
ip addr 172.16.5.5 255.255.255.0

int g0/0
ip addr 10.10.35.5 255.255.255.0
```

```
no shu
ip ospf 1 area 0
ip ospf message 1 md5 CCNP
ip ospf auth message

int g0/1
ip addr 10.10.45.5 255.255.255.0
no shu
ip ospf 1 area 1

router ospf 1
redistribute connected subnets
exi
```

upvoted 1 times

saiyuki1209 8 months, 1 week ago

why needs cost 10 on R5?

upvoted 1 times

HungarianDish 7 months ago

Yeah, I do not think either that cost is needed there.

upvoted 1 times

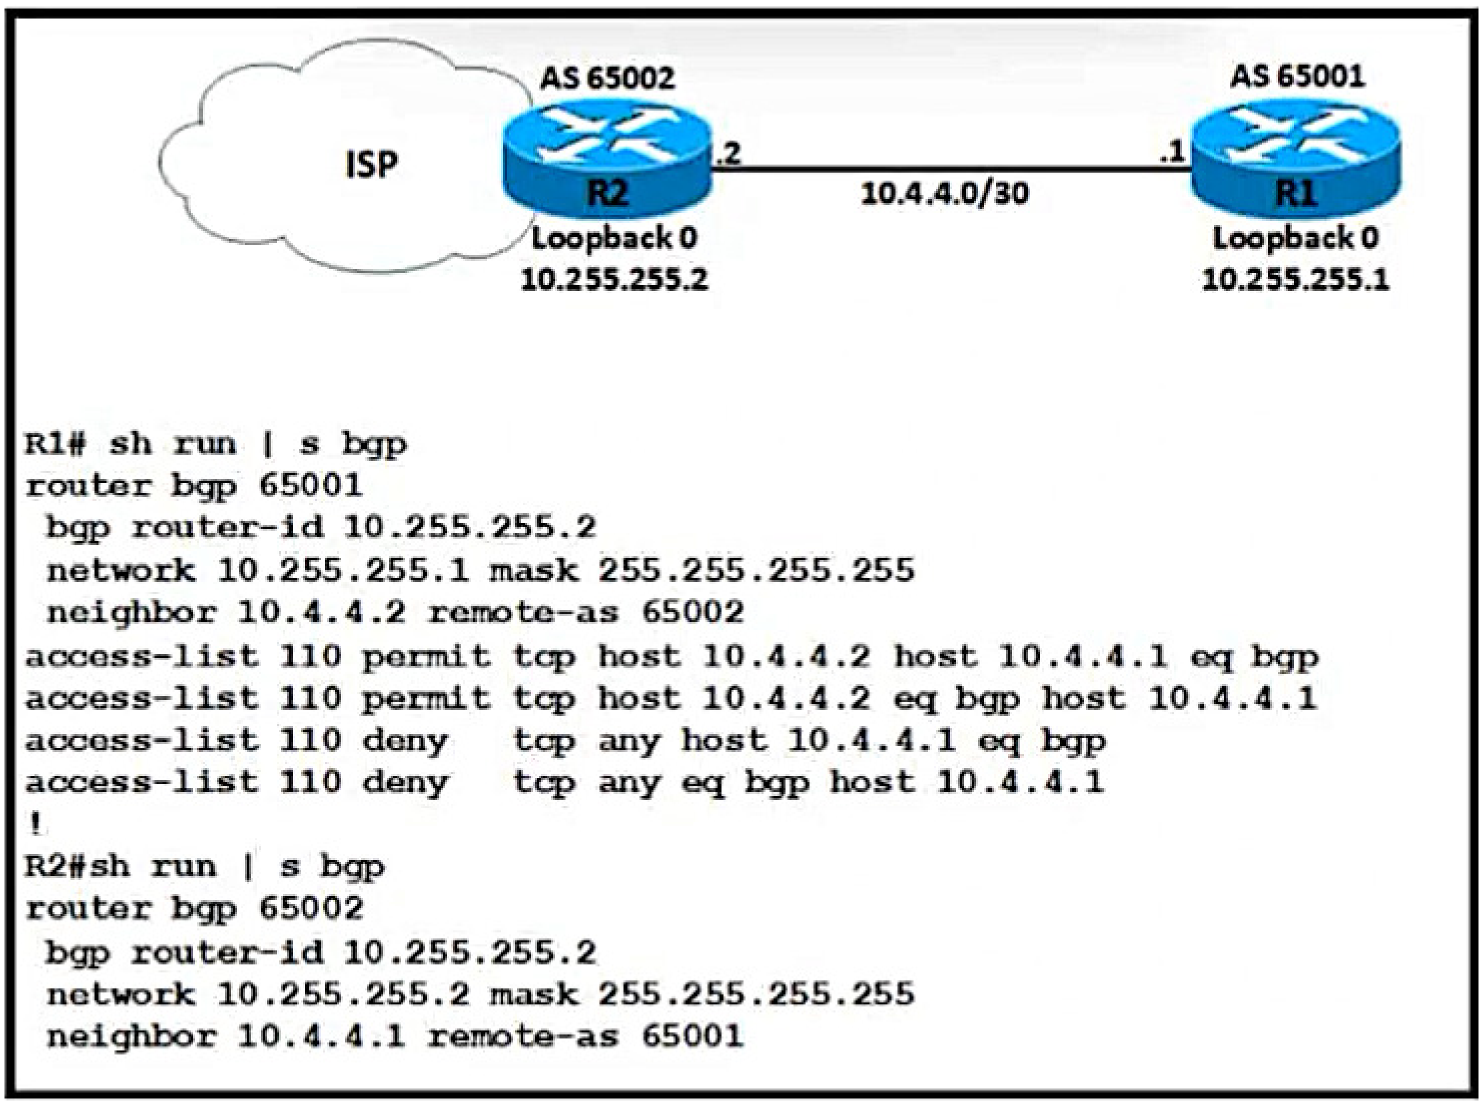

```
R1# sh run | s bgp
router bgp 65001
 bgp router-id 10.255.255.2
 network 10.255.255.1 mask 255.255.255.255
 neighbor 10.4.4.2 remote-as 65002
access-list 110 permit tcp host 10.4.4.2 host 10.4.4.1 eq bgp
access-list 110 permit tcp host 10.4.4.2 eq bgp host 10.4.4.1
access-list 110 deny   tcp any host 10.4.4.1 eq bgp
access-list 110 deny   tcp any eq bgp host 10.4.4.1
!
R2#sh run | s bgp
router bgp 65002
 bgp router-id 10.255.255.2
 network 10.255.255.2 mask 255.255.255.255
 neighbor 10.4.4.1 remote-as 65001
```

Refer to the exhibit. A network engineer notices that R1 and R2 cannot establish an eBGP peering. The following messages appear in the log:

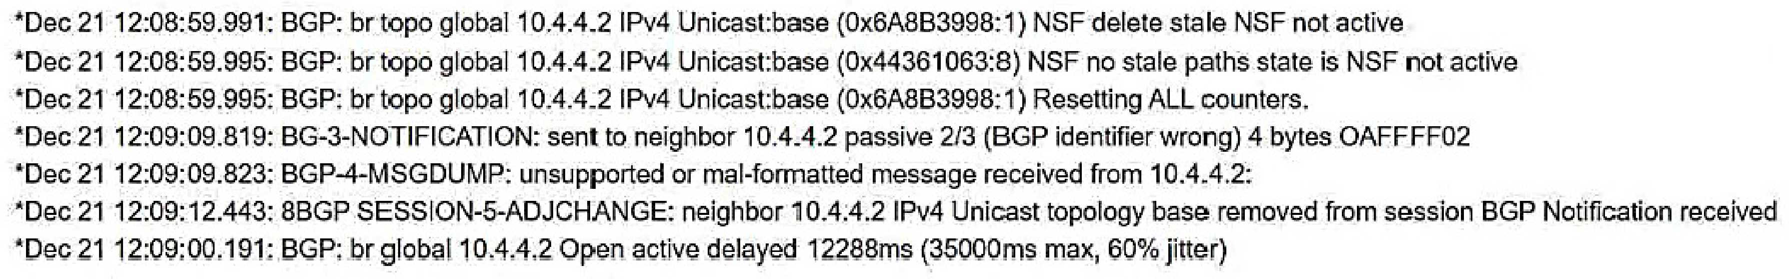

```
*Dec 21 12:08:59.991: BGP: br topo global 10.4.4.2 IPv4 Unicast:base (0x6A8B3998:1) NSF delete stale NSF not active
*Dec 21 12:08:59.995: BGP: br topo global 10.4.4.2 IPv4 Unicast:base (0x44361063:8) NSF no stale paths state is NSF not active
*Dec 21 12:08:59.995: BGP: br topo global 10.4.4.2 IPv4 Unicast:base (0x6A8B3998:1) Resetting ALL counters.
*Dec 21 12:09:09.819: BG-3-NOTIFICATION: sent to neighbor 10.4.4.2 passive 2/3 (BGP identifier wrong) 4 bytes OAFFFF02
*Dec 21 12:09:09.823: BGP-4-MSGDUMP: unsupported or mal-formatted message received from 10.4.4.2:
*Dec 21 12:09:12.443: 8BGP SESSION-5-ADJCHANGE: neighbor 10.4.4.2 IPv4 Unicast topology base removed from session BGP Notification received
*Dec 21 12:09:00.191: BGP: br global 10.4.4.2 Open active delayed 12288ms (35000ms max, 60% jitter)
```

Which configuration must the engineer apply to R1 to restore the eBGP peering?

A.
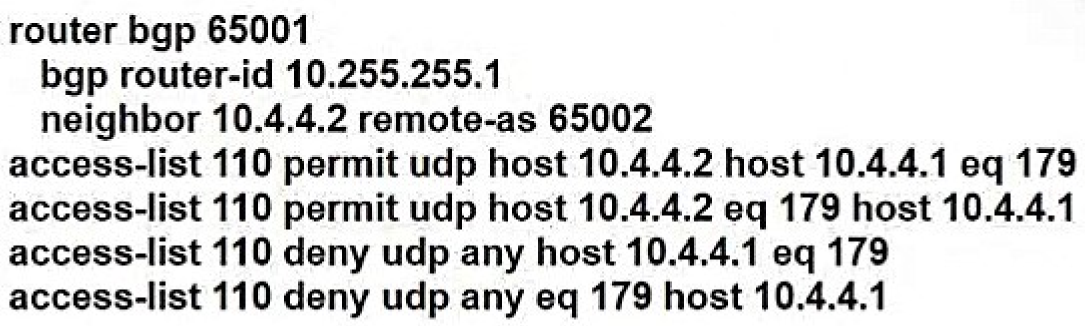

```
router bgp 65001
  bgp router-id 10.255.255.1
  neighbor 10.4.4.2 remote-as 65002
access-list 110 permit udp host 10.4.4.2 host 10.4.4.1 eq 179
access-list 110 permit udp host 10.4.4.2 eq 179 host 10.4.4.1
access-list 110 deny udp any host 10.4.4.1 eq 179
access-list 110 deny udp any eq 179 host 10.4.4.1
```

B.
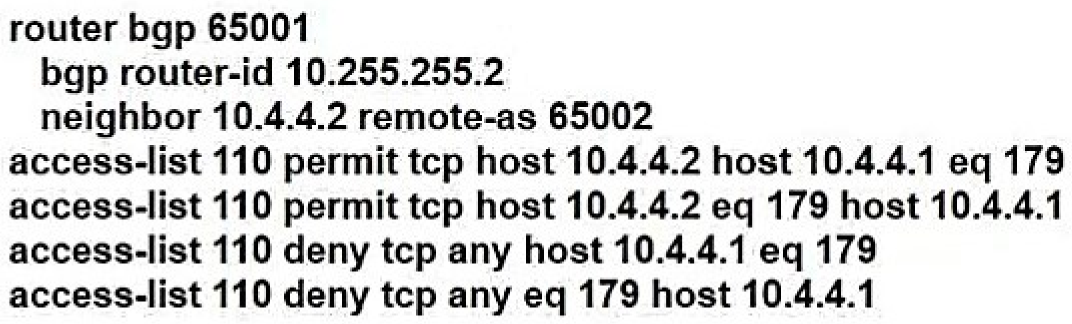

```
router bgp 65001
  bgp router-id 10.255.255.2
  neighbor 10.4.4.2 remote-as 65002
access-list 110 permit tcp host 10.4.4.2 host 10.4.4.1 eq 179
access-list 110 permit tcp host 10.4.4.2 eq 179 host 10.4.4.1
access-list 110 deny tcp any host 10.4.4.1 eq 179
access-list 110 deny tcp any eq 179 host 10.4.4.1
```

C.
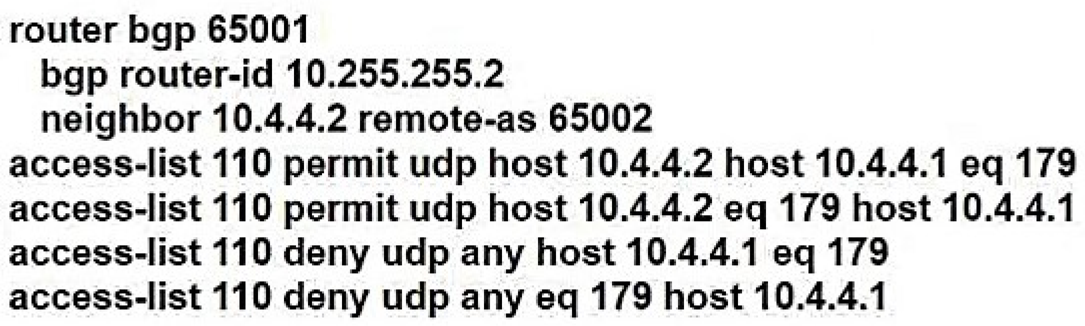

```
router bgp 65001
  bgp router-id 10.255.255.2
  neighbor 10.4.4.2 remote-as 65002
access-list 110 permit udp host 10.4.4.2 host 10.4.4.1 eq 179
access-list 110 permit udp host 10.4.4.2 eq 179 host 10.4.4.1
access-list 110 deny udp any host 10.4.4.1 eq 179
access-list 110 deny udp any eq 179 host 10.4.4.1
```

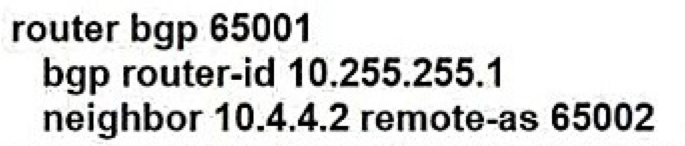

```
router bgp 65001
  bgp router-id 10.255.255.1
  neighbor 10.4.4.2 remote-as 65002
```

D.

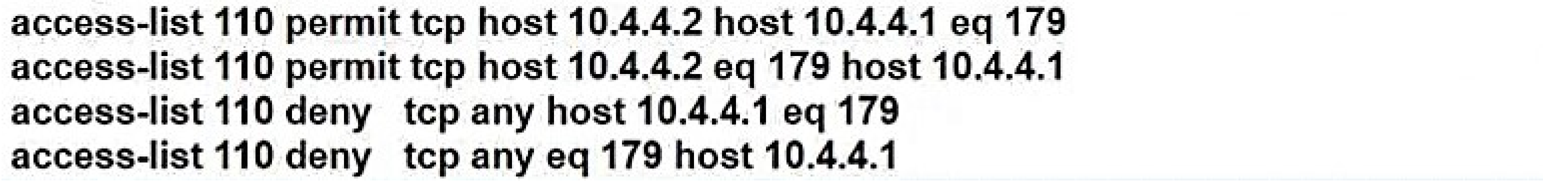

```
access-list 110 permit tcp host 10.4.4.2 host 10.4.4.1 eq 179
access-list 110 permit tcp host 10.4.4.2 eq 179 host 10.4.4.1
access-list 110 deny   tcp any host 10.4.4.1 eq 179
access-list 110 deny   tcp any eq 179 host 10.4.4.1
```

Correct Answer: *D*

782f5f0 4 months, 2 weeks ago

A

Router-id is false
UDP (transport protocol to use)

upvoted 1 times

robi1020 6 months, 1 week ago

Router-id and protocol (TCP) BGP uses TCP not UDP

upvoted 3 times

HungarianDish 6 months, 3 weeks ago

Selected Answer: D

Solution: Router-id on R1 under bgp process needs to be corrected.

upvoted 3 times

SIMULATION

-

Guidelines

-

This is a lab item in which tasks will be performed on virtual devices.

- Refer to the Tasks tab to view the tasks for this lab item.
- Refer to the Topology tab to access the device console(s) and perform the tasks.
- Console access is available for all required devices by clicking the device icon or using the tab(s) above the console window.
- All necessary preconfigurations have been applied.
- Do not change the enable password or hostname for any device.
- Do not replace existing routing policies or configurations.
- Save your configurations to NVRAM before moving to the next item.
- Click Next at the bottom of the screen to submit this lab and move to the next question.
- When Next is clicked, the lab closes and cannot be reopened.

Topology

-

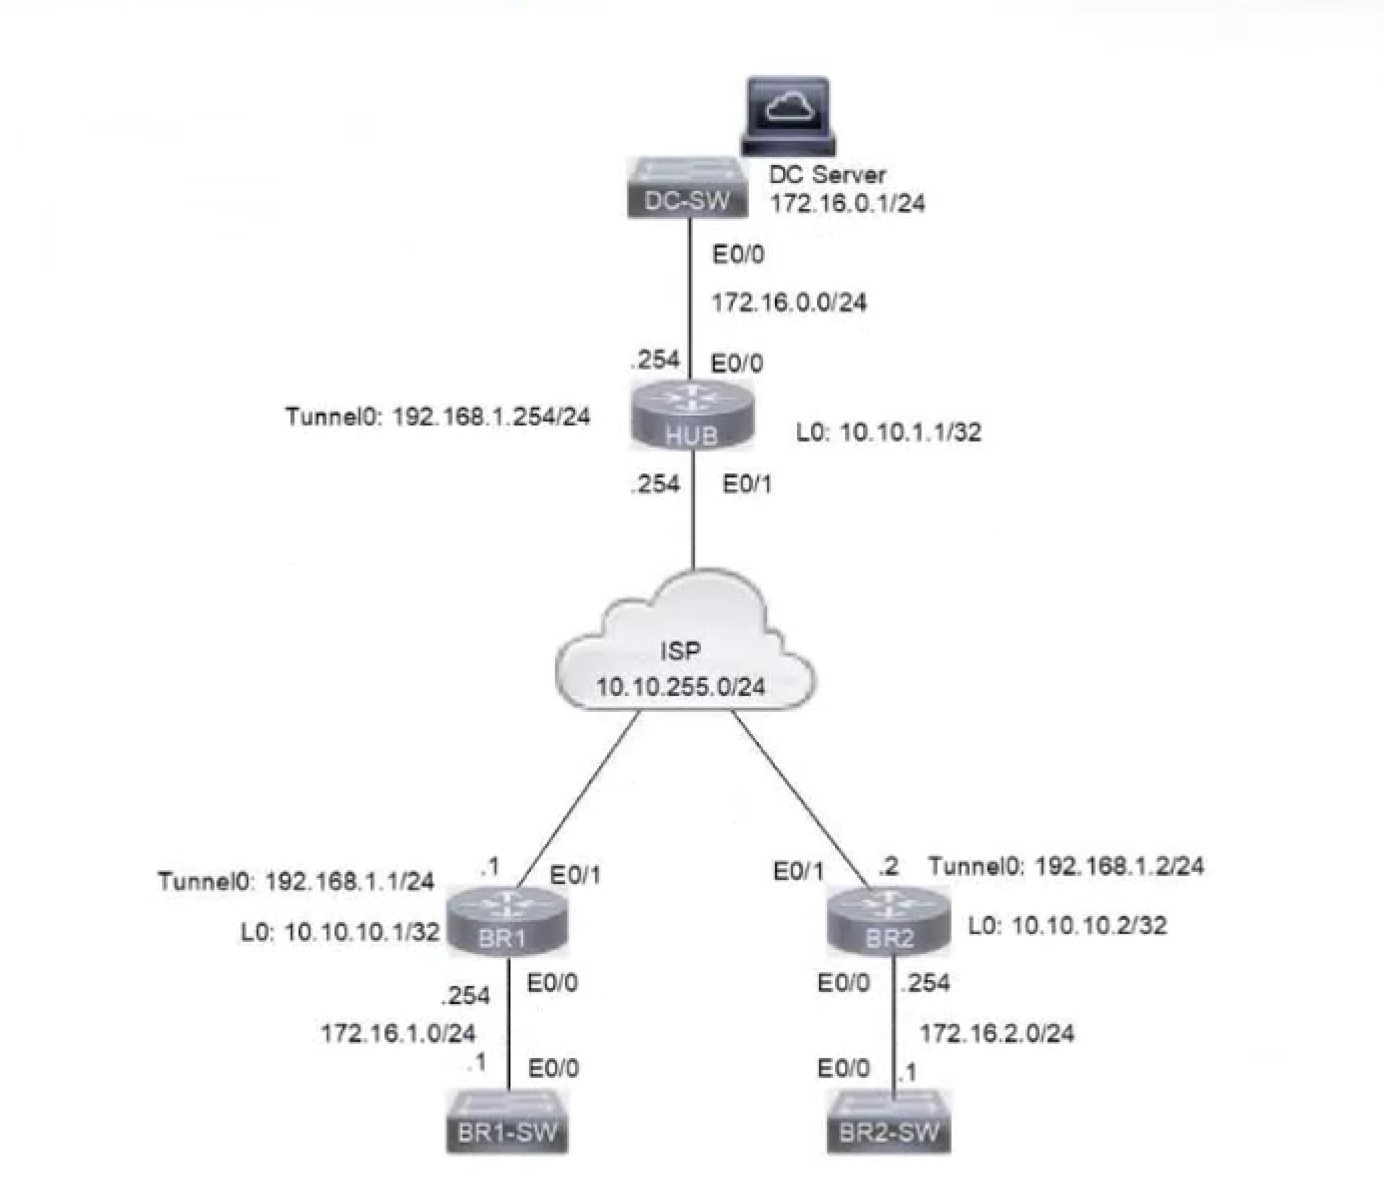

Topology Diagram

Tasks

-

A DMVPN network is preconfigured with tunnel 0 IP address 192.168.1.254 on the HUB, IP connectivity, crypto policies, profiles, and EIGRP AS

100. The NHRP password is ccnp123, and the network ID and tunnel key is EIGRP ASN. Do not introduce a static route. Configure DMVPN connectivity between routers BR1 and BR2 to the HUB router using physical interface as the tunnel source to achieve these goals:

1. Configure NHRP authentication, static IP-to-NBMA address maps, hold time 5 minutes, network ID, and server on branch router BR1.
2. Configure NHRP authentication, static IP-to-NBMA address maps, hold time 5 minutes, network ID, and server on branch router BR2.
3. Ensure that packet fragmentation is done before encryption to account for GRE and IPsec header and allow a maximum TCP segment size of 1360 on an IP MTU of 1400 on the tunnel interfaces of both branch routers.
4. Apply an IPsec profile to the tunnel. Verify that direct spoke-to-spoke tunnel is functional between branch routers BR1 and BR2 by using traceroute to Ethernet 0/0 IP address to get a full score.

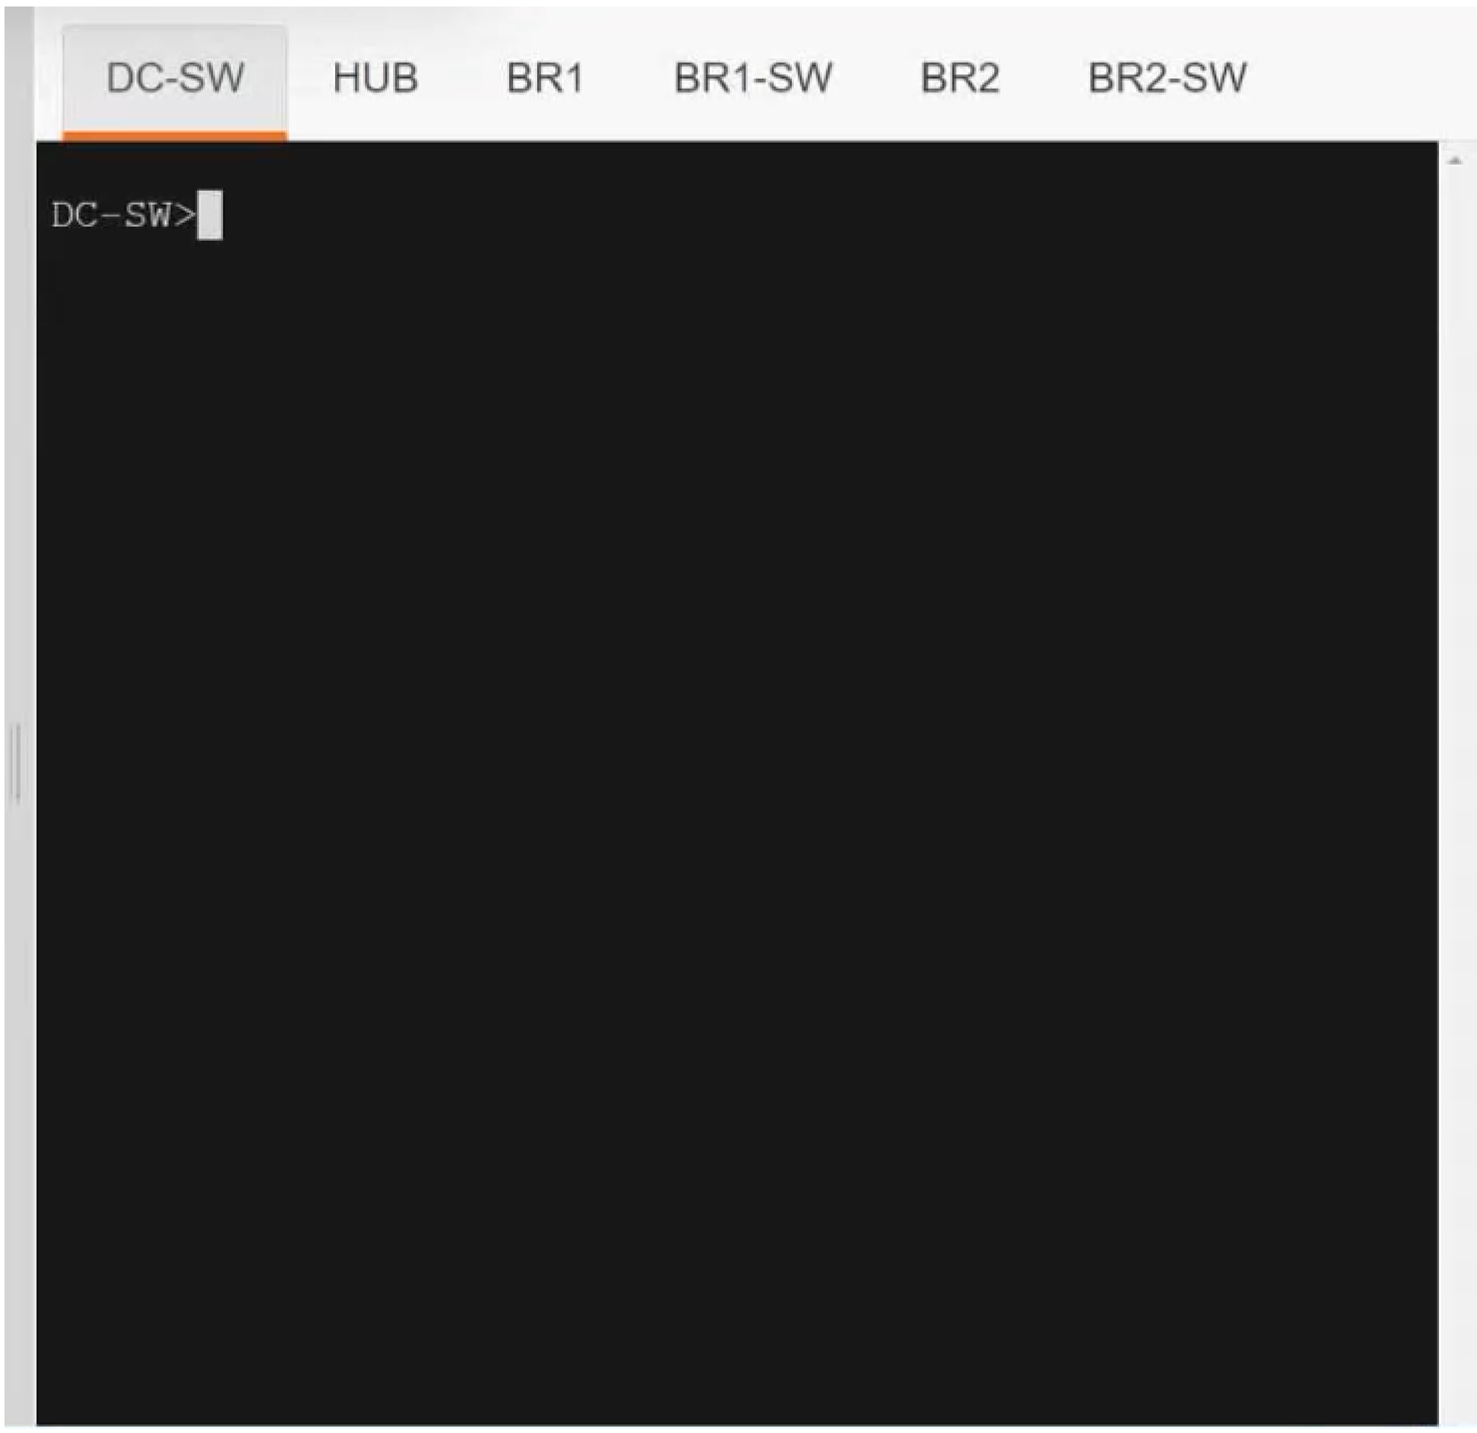

```
DC-SW>
```

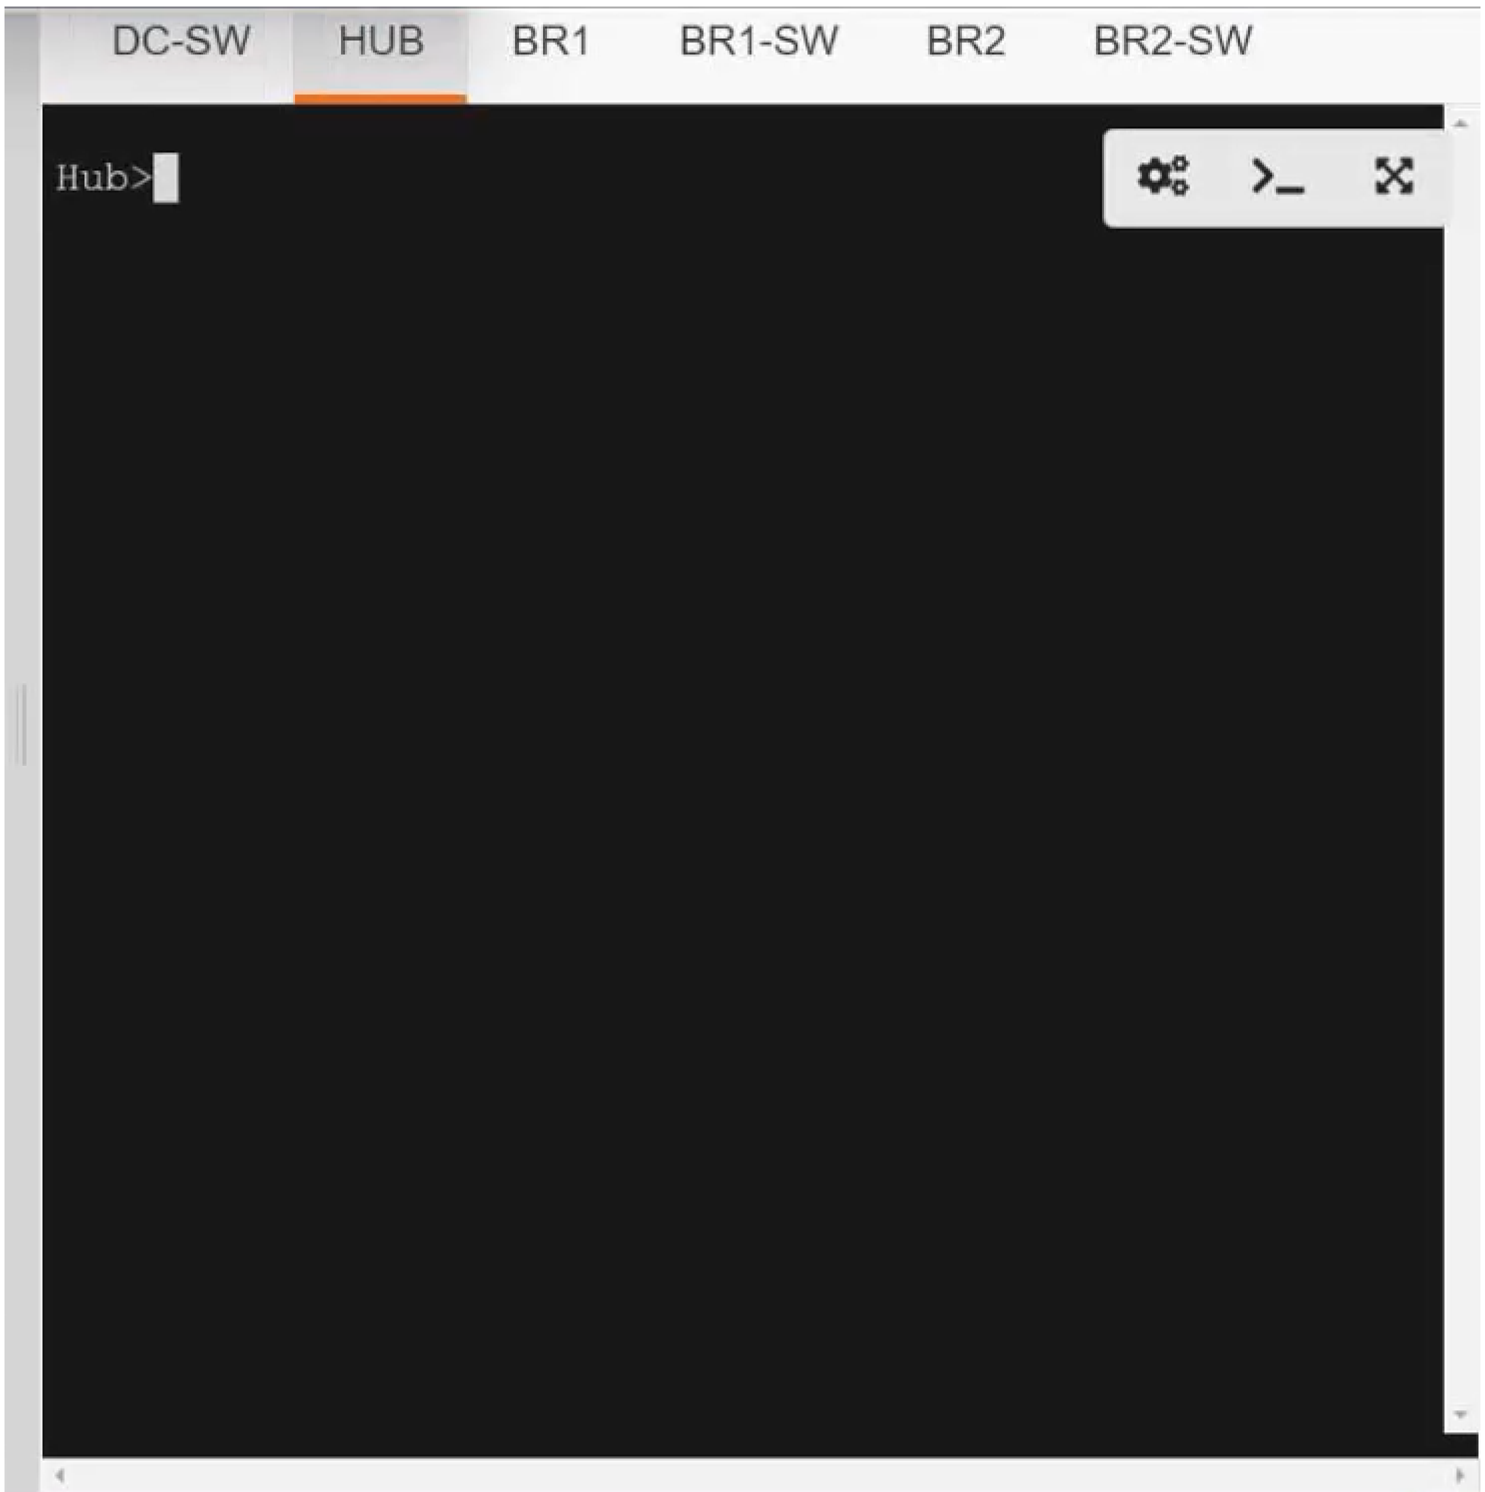

```
Hub>
```

```
BR1>
```

```
BR1-SW con0 is now available



Press RETURN to get started.
```

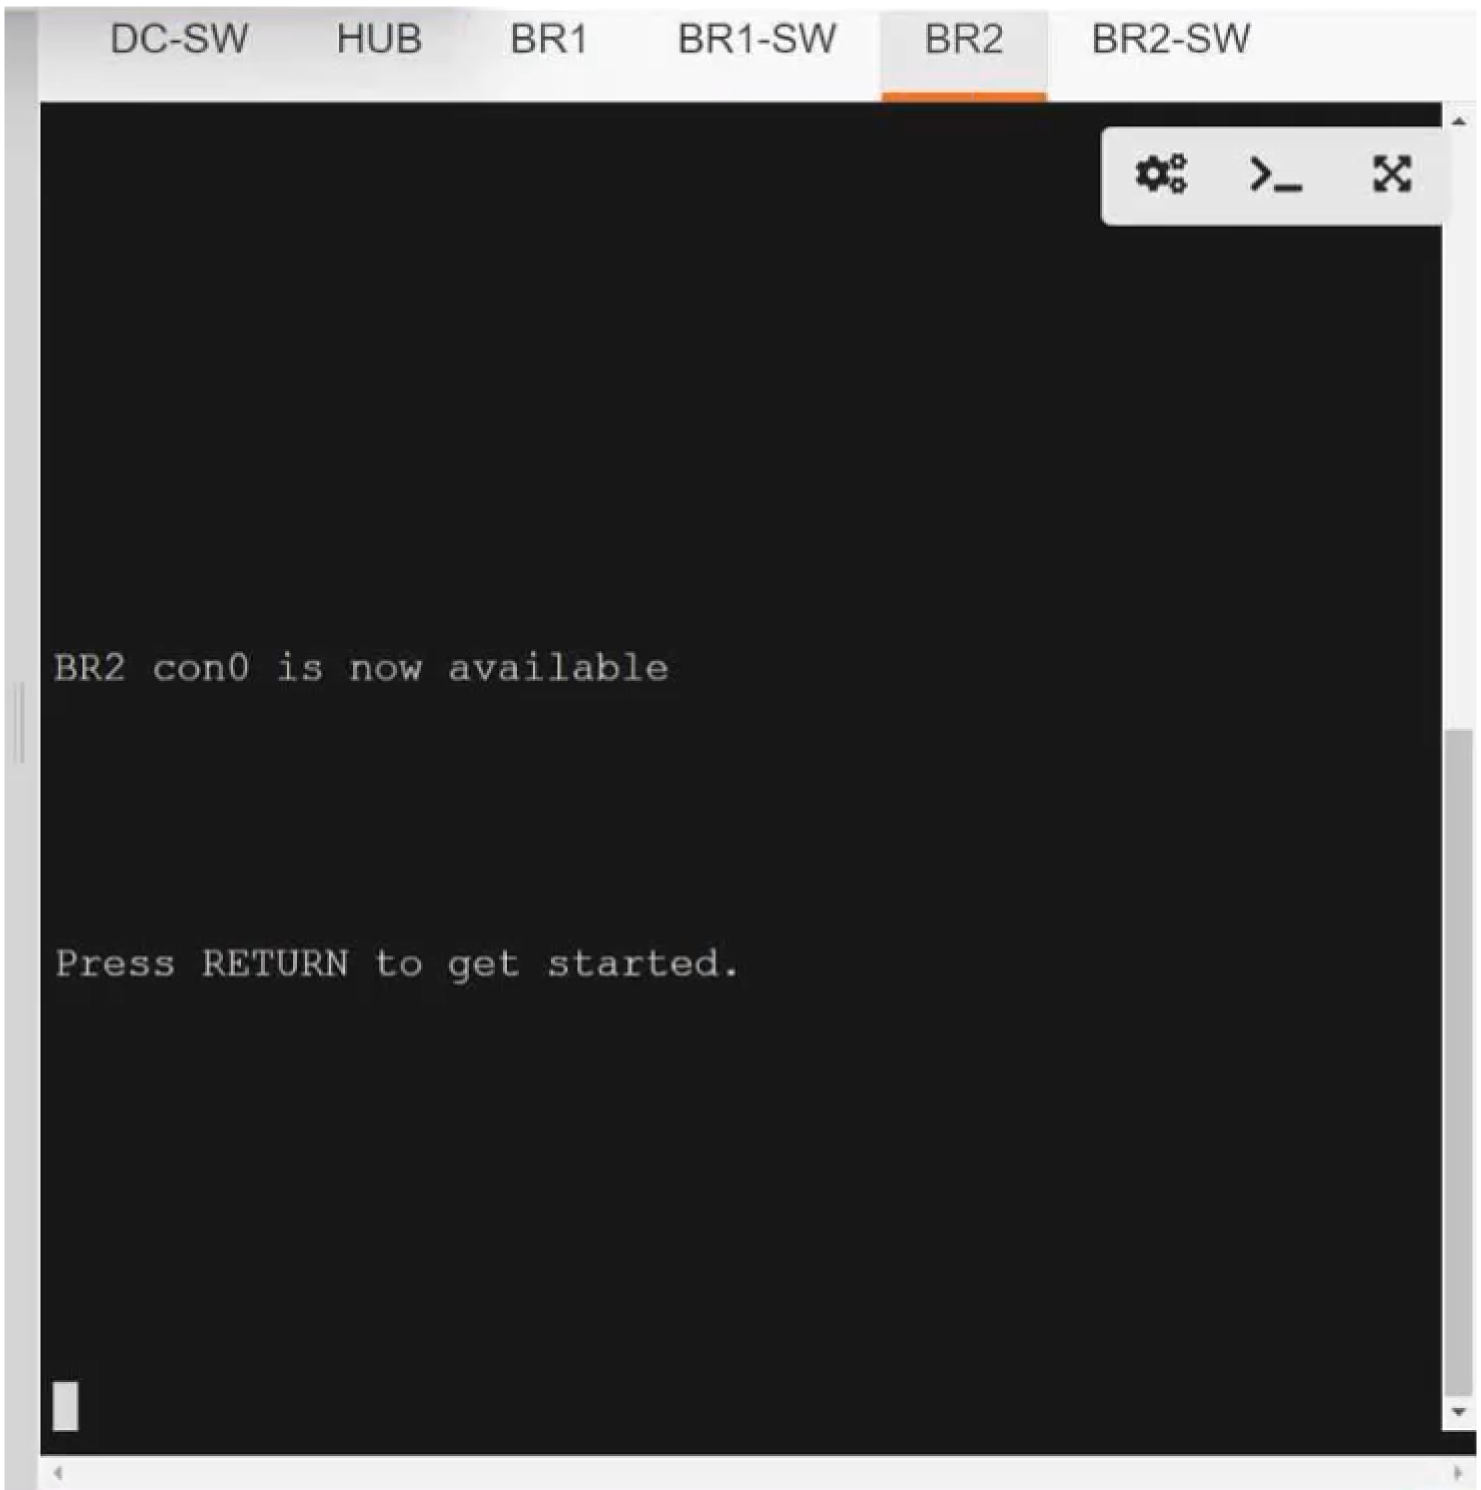

```
BR2 con0 is now available



Press RETURN to get started.
```

```
BR1:

Interface Tunnel0
 IP address 192.168.1.1 255.255.255.0
no ip redirects
ip mtu 1400
ip nhrp authentication ccnp123
ip nhrp map 192.168.1.254 10.10.255.254
ip nhrp map multicast 10.10.255.254
ip nhrp network-id 100
ip nhrp holdtime 5
ip nhrp nhs 192.168.1.254
ip nhrp shortcut
ip tcp adjust-mss 1360
delay 1000
tunnel source 10.10.255.1
tunnel destination 10.10.255.254
tunnel key 100
```

Correct Answer:

```
BR2:

Interface Tunnel0
 IP address 192.168.1.2 255.255.255.0
no ip redirects
ip mtu 1400
ip nhrp authentication ccnp123
ip nhrp map 192.168.1.254 10.10.255.254
ip nhrp map multicast 10.10.255.254
ip nhrp network-id 100
ip nhrp holdtime 5
ip nhrp nhs 192.168.1.254
ip nhrp shortcut
ip tcp adjust-mss 1360
delay 1000
tunnel source 10.10.10.2
tunnel destination 10.10.255.254
tunnel key 100
```

⊟ **Brand** 3 months, 2 weeks ago

I got this exact sim in the exam and I'm here to confirm that presented answer is almost correct. Just make sure you use the lines below:
ip nhrp holdtime 300 (you define the holdtime using seconds)
tunnel mode gre multipoint (as the sim asks you to confirm direct communication between Spokes. Static destination is not what you looking for here)

upvoted 2 times

⊟ **Brand** 3 months, 2 weeks ago

Also, don't forget to add following lines in the global configuration mode as HungarianDish explained.

crypto ipsec df-bit clear
crypto ipsec fragmentation before-encrypt

Because the sim asks you to make sure fragmentation is being done before the encryption with IPSec.

upvoted 3 times

⊟ **HungarianDish** 6 months, 1 week ago

Sim was the same as here. Did not get full score for this sim on real exam, and could not figure out, what was missing. Still passed the exam. New question: MP-BGP NLRI attributes: RD, IPv4 Prefix, Next Hop, VPN Label
https://networklessons.com/mpls/mpls-layer-3-vpn-explained

upvoted 4 times

⊟ **keesu** 5 months, 4 weeks ago

congratz on passing the exam!
Thank you for your valuable comments throughout the Qs!

upvoted 3 times

⊟ **ParisaAlipoor** 5 months, 2 weeks ago

Thank you for your valuable comments... can you please describe the new question?
MP-BGP NLRI attributes: RD, IPv4 Prefix, Next Hop, VPN Label

upvoted 1 times

⊟ **HungarianDish** 5 months ago

List the MP-BGP NLRI attributes. Answer: RD, IPv4 Prefix, Next Hop, VPN Label

upvoted 1 times

GReddy2323 6 months, 4 weeks ago

I would like to lab this in CML, but what type of basic configuration does the ISP router need? I always struggle whenever an "ISP" router is put in the topology because I don't know what basic configuration it needs to get it to work.

upvoted 1 times

cir_ 6 months, 3 weeks ago

You can use a layer 2 switch with no config to replicate the connectivity in this scenario

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

L2 switch in the middle is the best! I saw also different topology in labs (e.g. boson netsim), where a router sits in the middle as ISP, and then static routes are added an all routers to reach each others WAN IPs. It is unnecessary to have such a topology for practicing DMVPN. L2 switch is perfect.

upvoted 1 times

HungarianDish 7 months ago

```
hostname HUB
!
crypto isakmp policy 5
hash md5
authentication pre-share
crypto isakmp key cisco address 0.0.0.0
!
crypto ipsec transform-set ciscoset esp-3des
mode tunnel
!
crypto ipsec profile ciscoprofile
set transform-set ciscoset
!
interface Tunnel0
ip address 192.168.1.254 255.255.255.0
no ip redirects
no ip next-hop-self eigrp 100
no ip split-horizon eigrp 100
ip nhrp authentication ccnp123
ip nhrp network-id 1
ip nhrp holdtime 300
ip nhrp redirect
tunnel source GigabitEthernet0/1
tunnel mode gre multipoint
tunnel key 100
tunnel protection ipsec profile ciscoprofile
!
router eigrp 100
network 10.10.1.1 0.0.0.0
network 192.168.1.0
!
crypto ipsec df-bit clear
crypto ipsec fragmentation before-encrypt
```

upvoted 1 times

chris110 3 months, 1 week ago

Need to configure the HUB as well?

upvoted 1 times

Almylle 5 months, 3 weeks ago

No tunnel destination ?

upvoted 1 times

chris110 3 months, 1 week ago

i dont think we need because of tunnel mode gre multipoint

upvoted 1 times

HungarianDish 7 months ago

```
int tu 0
ip nhrp map multicast dynamic
```

(left out from previous comment by mistake)

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Some correction for transform-set:
https://community.cisco.com/t5/vpn/dmvpn-tunnel-versus-transport-mode/td-p/1544252
Transport mode actually is recommended mode for DMVPN, because it saves 20 bytes overhead.

```
crypto ipsec transform-set ciscoset esp-3des
mode transport
```

upvoted 1 times

HungarianDish 6 months, 2 weeks ago

This document states that we would need "mode tunnel" if we want to do pre-fragmentation. Probably we can leave it in the default mode, which is tunnel.
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_dplane/configuration/xe-16-10/sec-ipsec-data-plane-xe-16-10-book/sec-pre-frag-vpns.html
Pre-fragmentation for IPsec VPNs operates in IPsec tunnel mode and IPsec tunnel mode with GRE, but not with IPsec transport mode.

upvoted 1 times

HungarianDish 7 months ago

```
hostname BR1
!
interface Tunnel0
ip address 192.168.1.1 255.255.255.0
no ip redirects
ip mtu 1400
ip nhrp authentication ccnp123
ip nhrp map multicast 10.10.255.254
ip nhrp map 192.168.1.254 10.10.255.254
ip nhrp network-id 1
ip nhrp holdtime 300
ip nhrp nhs 192.168.1.254
ip tcp adjust-mss 1360
tunnel source GigabitEthernet0/1
tunnel mode gre multipoint
tunnel key 100
tunnel protection ipsec profile ciscoprofile
!
router eigrp 100
network 10.10.10.1 0.0.0.0
network 192.168.1.0
!
crypto ipsec df-bit clear
crypto ipsec fragmentation before-encrypt
```

upvoted 1 times

HungarianDish 7 months ago

```
interface Tunnel0
ip nhrp shortcut
```

(left out from previous comment by mistake)

upvoted 1 times

HungarianDish 7 months ago

```
hostname BR2
!
interface Tunnel0
ip address 192.168.1.2 255.255.255.0
no ip redirects
ip mtu 1400
ip nhrp authentication ccnp123
ip nhrp map multicast 10.10.255.254
ip nhrp map 192.168.1.254 10.10.255.254
ip nhrp network-id 1
ip nhrp holdtime 300
ip nhrp nhs 192.168.1.254
ip tcp adjust-mss 1360
tunnel source GigabitEthernet0/1
tunnel mode gre multipoint
tunnel key 100
tunnel protection ipsec profile ciscoprofile
!
router eigrp 100
network 10.10.10.2 0.0.0.0
network 192.168.1.0
!
crypto ipsec df-bit clear
crypto ipsec fragmentation before-encrypt
```

upvoted 1 times

HungarianDish 7 months ago

```
interface Tunnel0
ip nhrp shortcut
```

(left out from previous comment by mistake)

upvoted 1 times

HungarianDish 7 months ago

3. Ensure that packet fragmentation is done before encryption
=> For me, it suggests that following config is required (global config mode or interface config):

```
crypto ipsec df-bit clear
crypto ipsec fragmentation before-encrypt
```

upvoted 3 times

HungarianDish 7 months ago

https://networkengineering.stackexchange.com/questions/11283/pre-fragmentation-for-ipsec-vpns-on-cisco-routers

I added it in global config mode.

upvoted 1 times

HungarianDish 6 months, 3 weeks ago

Interestingly, "crypto ipsec fragmentation before-encrypt" is never displayed in the running config. Maybe it is due to this bug:

de gre

Question #474 *Topic 1*

What is the downstream unsolicited distribution method in MPLS?

A. It advertises labels to peers only when the peer requests.

B. It sends a unicast hello message to a specific LSR.

C. It sends a unicast hello message to a specific LER.

D. It advertises labels to peers without peer request.

Correct Answer: *D*

robi1020 5 months ago

Selected Answer: D

Unsolicited Downstream :
The MPLS architecture also allows an LSR to distribute bindings to LSRs that have not explicitly requested them. This is known as "unsolicited downstream" label distribution. This method is used in LDP & BGP-LU(RFC 3107)

upvoted 2 times

Question #475 *Topic 1*

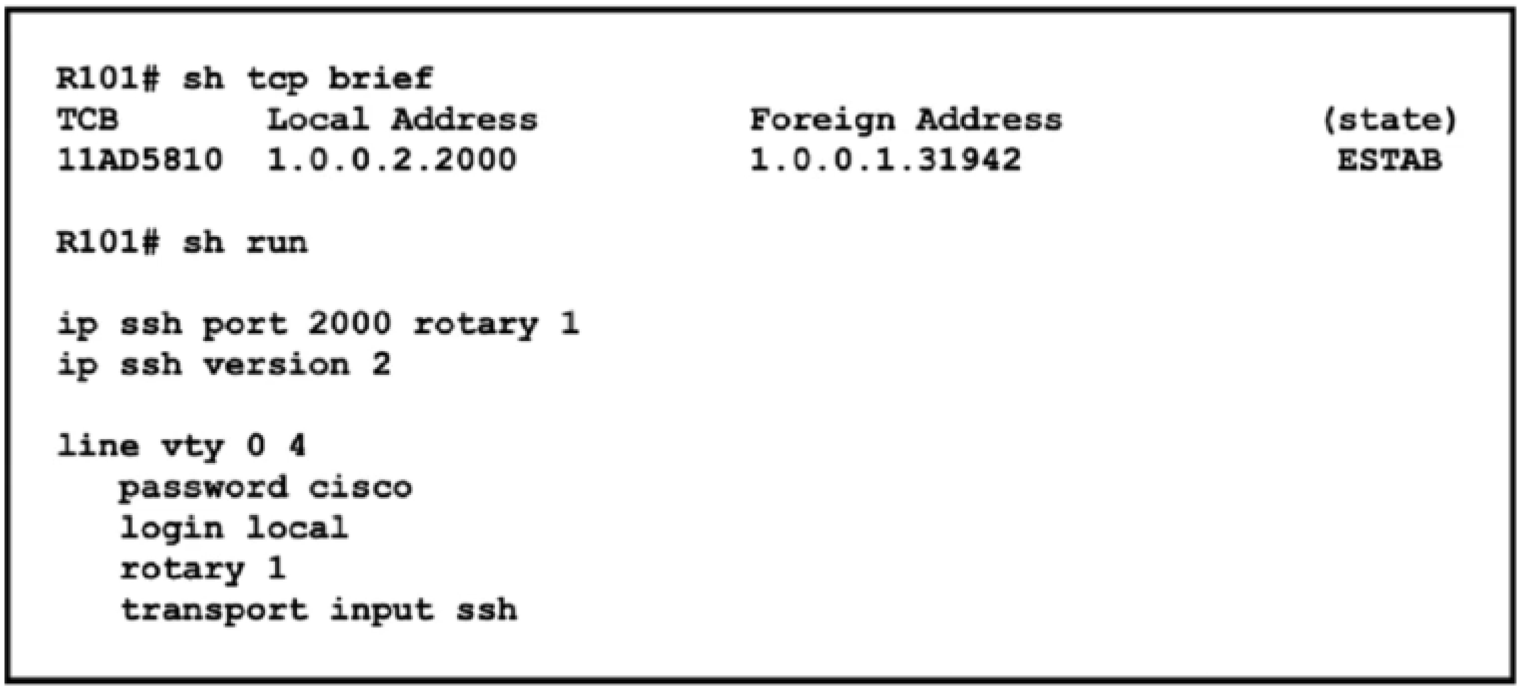

```
R101# sh tcp brief
TCB       Local Address           Foreign Address              (state)
11AD5810  1.0.0.2.2000            1.0.0.1.31942                 ESTAB

R101# sh run

ip ssh port 2000 rotary 1
ip ssh version 2

line vty 0 4
   password cisco
   login local
   rotary 1
   transport input ssh
```

Refer to the exhibit. An engineer must configure router R101 for SSH access on ports 2001 through 2011. After the configuration, some expected ports were inaccessible. Which command resolves the issue?

A. ip ssh port 2001 rotary 11
line vty 0 4
transport input telnet

B. ip ssh port 2000 rotary 11
line vty 0 4
transport input ssh

C. ip ssh port 2000 rotary 1 11
line vty 0 4
transport input all

D. ip ssh port 2001 rotary 1 11
line vty 0 4
transport input ssh

Correct Answer: *D*

Tedmus 1 month, 1 week ago

Selected Answer: D

D is correct:

Link:
https://www.cisco.com/c/en/us/support/docs/security-vpn/secure-shell-ssh/212142-Configure-SSH-on-Tty-Lines-with-Menu-Opt.html
upvoted 1 times

Rman0059 1 month, 2 weeks ago

Selected Answer: D

D is correct
upvoted 1 times

Refer to the exhibit.

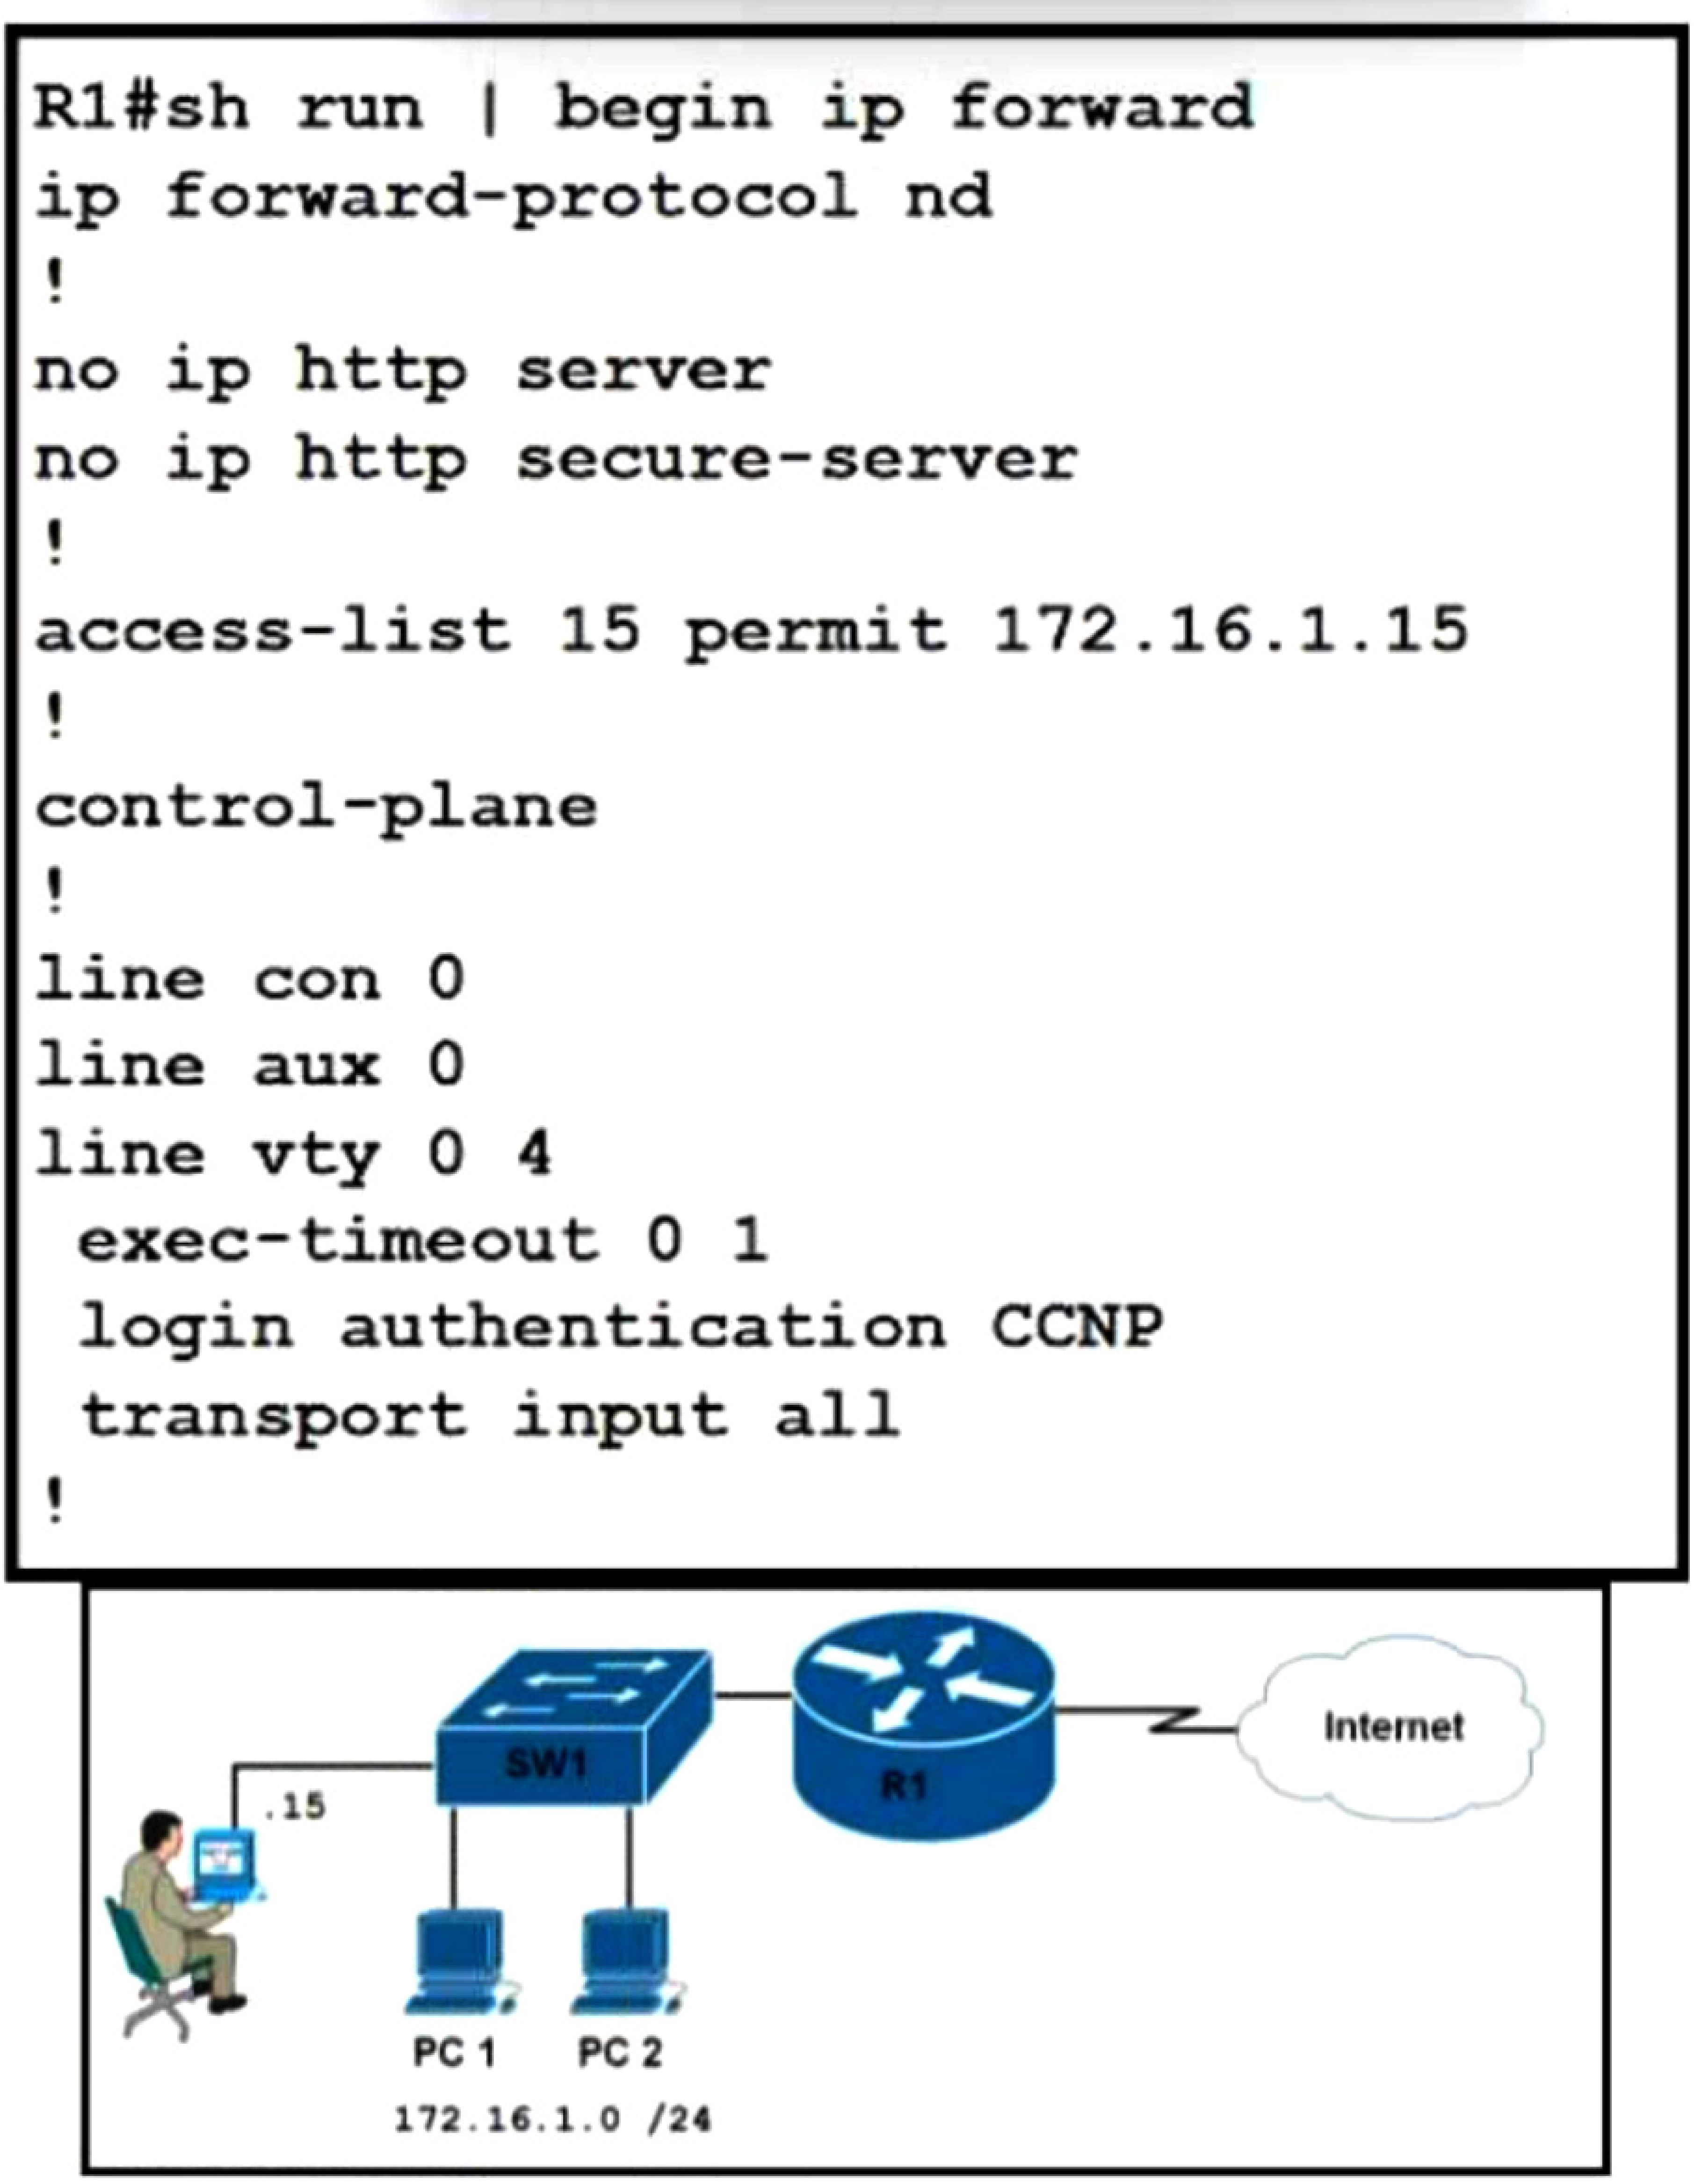

```
R1#sh run | begin ip forward
ip forward-protocol nd
!
no ip http server
no ip http secure-server
!
access-list 15 permit 172.16.1.15
!
control-plane
!
line con 0
line aux 0
line vty 0 4
 exec-timeout 0 1
 login authentication CCNP
 transport input all
!
```

After a RADIUS server fails AAA authentication, an engineer is trying to reestablish console access to a switch using the local password.

Which configuration reestablishes the console access to switch SW1 via AAA?

A.

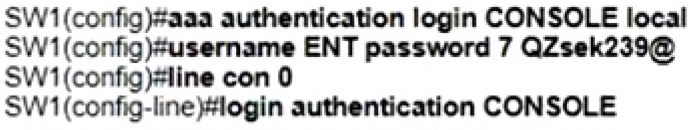

```
SW1(config)#aaa authentication login CONSOLE local
SW1(config)#username ENT password 7 QZsek239@
SW1(config)#line con 0
SW1(config-line)#login authentication CONSOLE
```

B.

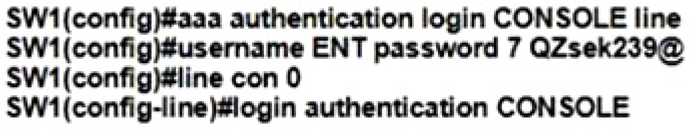

```
SW1(config)#aaa authentication login CONSOLE line
SW1(config)#username ENT password 7 QZsek239@
SW1(config)#line con 0
SW1(config-line)#login authentication CONSOLE
```

C.

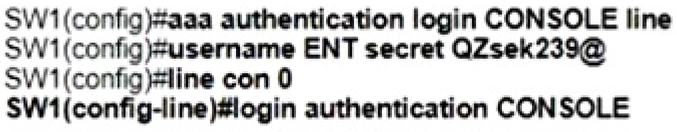

```
SW1(config)#aaa authentication login CONSOLE line
SW1(config)#username ENT secret QZsek239@
SW1(config)#line con 0
SW1(config-line)#login authentication CONSOLE
```

D.

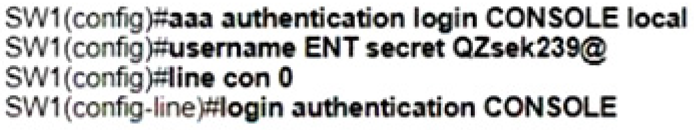

```
SW1(config)#aaa authentication login CONSOLE local
SW1(config)#username ENT secret QZsek239@
SW1(config)#line con 0
SW1(config-line)#login authentication CONSOLE
```

Correct Answer: *A*

⊟ **guy276465281819372** 4 months ago

Selected Answer: D

Password 7 expects an ENCRYPTED string after the command, you cannot just type the password. this is usually used to copy configuration from one switch to antoher. the only valid answer is D.

upvoted 3 times

⊟ **gpaulino** 4 months ago

Selected Answer: D

R5(config-router)#metric weights ?
<0-8> Type (Only TOS 0 supported)

R5(config-router)#metric weights 0 ?
<0-255> K1

R5(config-router)#metric weights 0 1 ?
<0-255> K2

R5(config-router)#metric weights 0 1 0 ?
<0-255> K3

R5(config-router)#metric weights 0 1 0 1 ?
<0-255> K4

R5(config-router)#metric weights 0 1 0 1 0 ?
<0-255> K5

upvoted 1 times

⊟ **Brand** 3 months, 2 weeks ago

dude...

upvoted 1 times

⊟ **Muste** 4 months, 1 week ago

Selected Answer: D

the correct anser is D

upvoted 2 times

⊟ **JJH3003** 4 months, 2 weeks ago

SW1(config)#username ENT password 7 QZsek239@
Invalid encrypted password: QZsek239@

Please set a password for username
SW1(config)#
==

"username ENT password 7" is expecting the encrypted password string to follow. The password provided in answer A is the unencrypted string which is invalid.

The correct command for the given password in answer a would be:
SW1(config)#username ENT password 7 0810765D0C1257444B2B

D is correct.

upvoted 2 times

⊟ **robi1020** 4 months, 2 weeks ago

Selected Answer: D

Not sure but D looks more correct. Because its "secret" ?

upvoted 2 times

⊟ **inteldarvid** 4 months, 2 weeks ago

Yes you have a reason, the answer correct is D, stronger password use username secret

upvoted 1 times

Question #477 Topic 1

Refer to the exhibit.

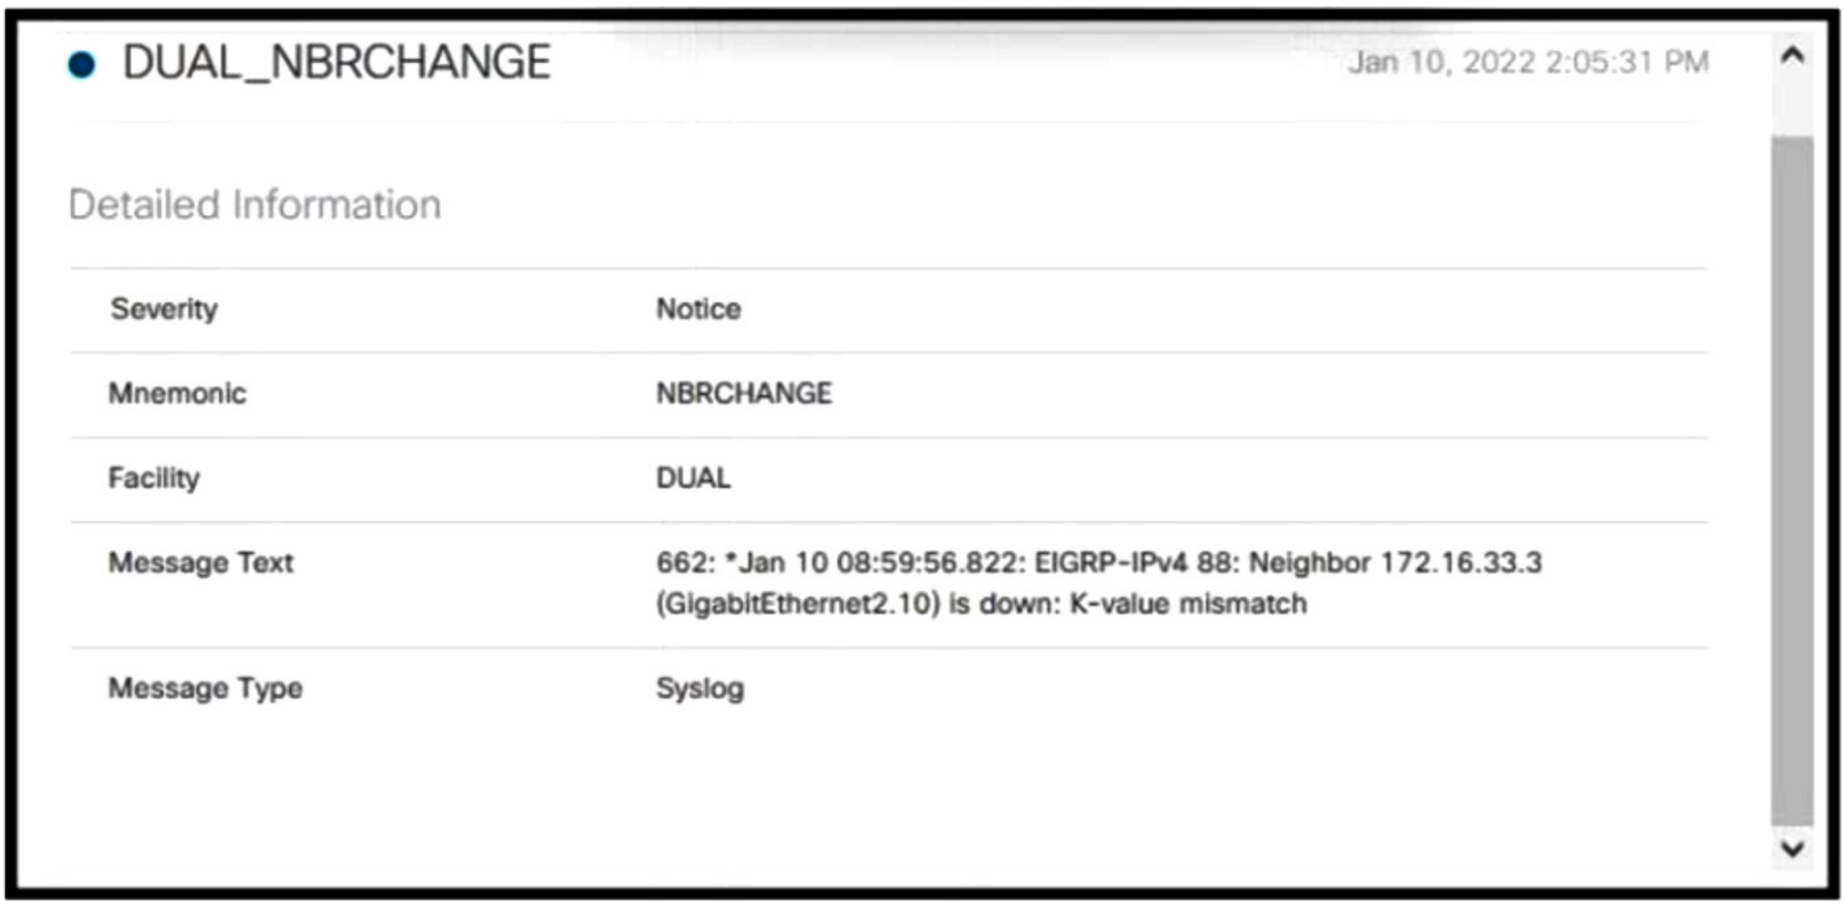

● DUAL_NBRCHANGE Jan 10, 2022 2:05:31 PM

Detailed Information

Severity	Notice
Mnemonic	NBRCHANGE
Facility	DUAL
Message Text	662: *Jan 10 08:59:56.822: EIGRP-IPv4 88: Neighbor 172.16.33.3 (GigabitEthernet2.10) is down: K-value mismatch
Message Type	Syslog

EIGRP peering was lost.

Which configuration resolves the issue?

A. 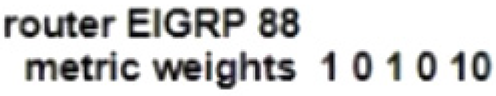

```
router EIGRP 88
 metric weights  1 0 1 0 10
```

B. 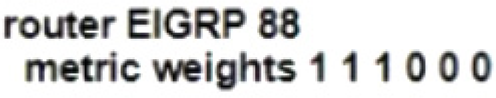

```
router EIGRP 88
 metric weights 1 1 1 0 0 0
```

C. 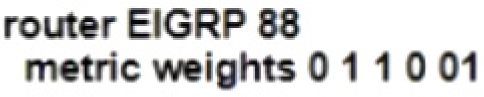

```
router EIGRP 88
 metric weights 0 1 1 0 01
```

D. 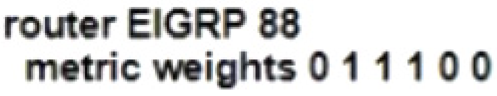

```
router EIGRP 88
 metric weights 0 1 1 1 0 0
```

Correct Answer: *A*

⊟ **ridonak230** 3 months ago

Selected Answer: D

Answer is "D" ...
I agree that K values are 5 , but when using "metric weights" command, the FIRST value is the TOS (Type Of Service) byte but as you can see it only supports a value of 0. The next values are for the actual K values:

```
R1(config)#router eigrp 88
R1(config-router)#metric weights ?
<0-8> Type Of Service (Only TOS 0 supported). <<<<<<<<
R1(config-router)#metric weights 0 ?
<0-255> K1
```

upvoted 2 times

⊟ **Muste** 4 months, 1 week ago

Selected Answer: D

the default metric should be like this the first field is TOS and it's mostly 0 the second field is K1 it should be 1 and the 4th field is K3 and it should be 1
R1(config-router)#metric weights 0 1 0 1 0 0

upvoted 2 times

⊟ **robi1020** 4 months, 2 weeks ago

Selected Answer: A

I think its A too. https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/15-s/ire-15-s-book/ire-wid-met.html

upvoted 1 times

Refer to the exhibit.

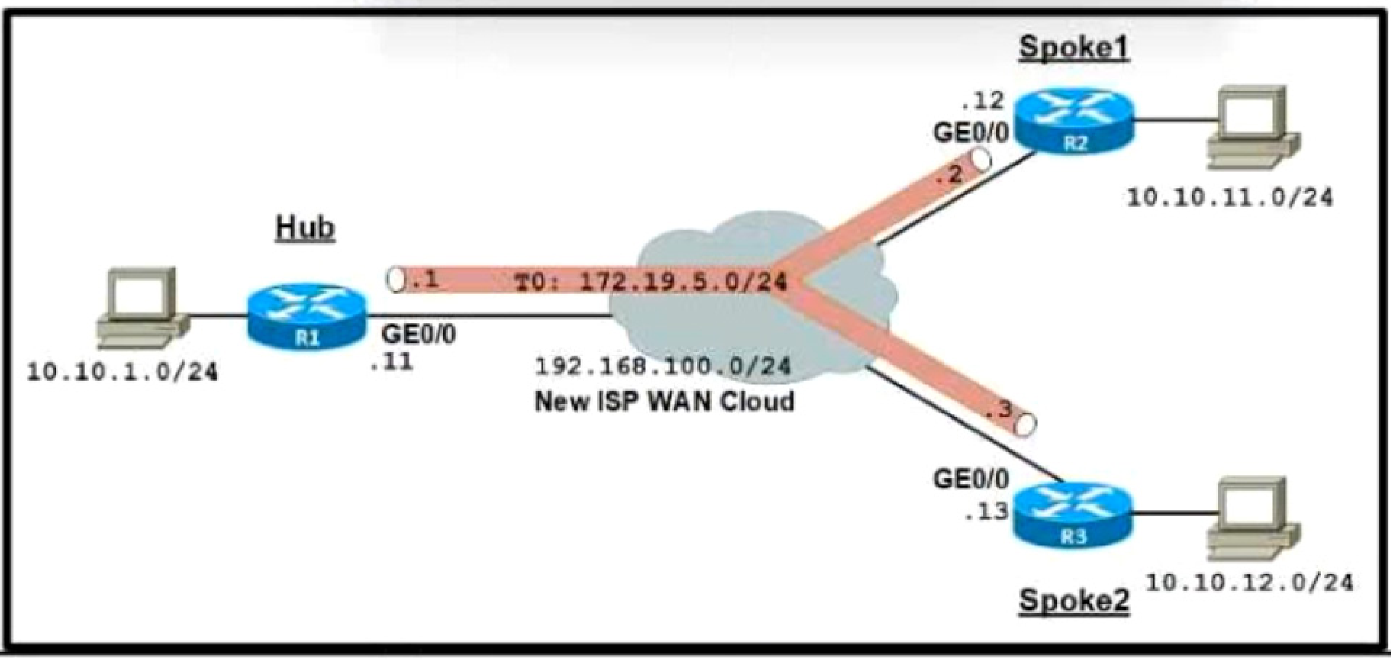

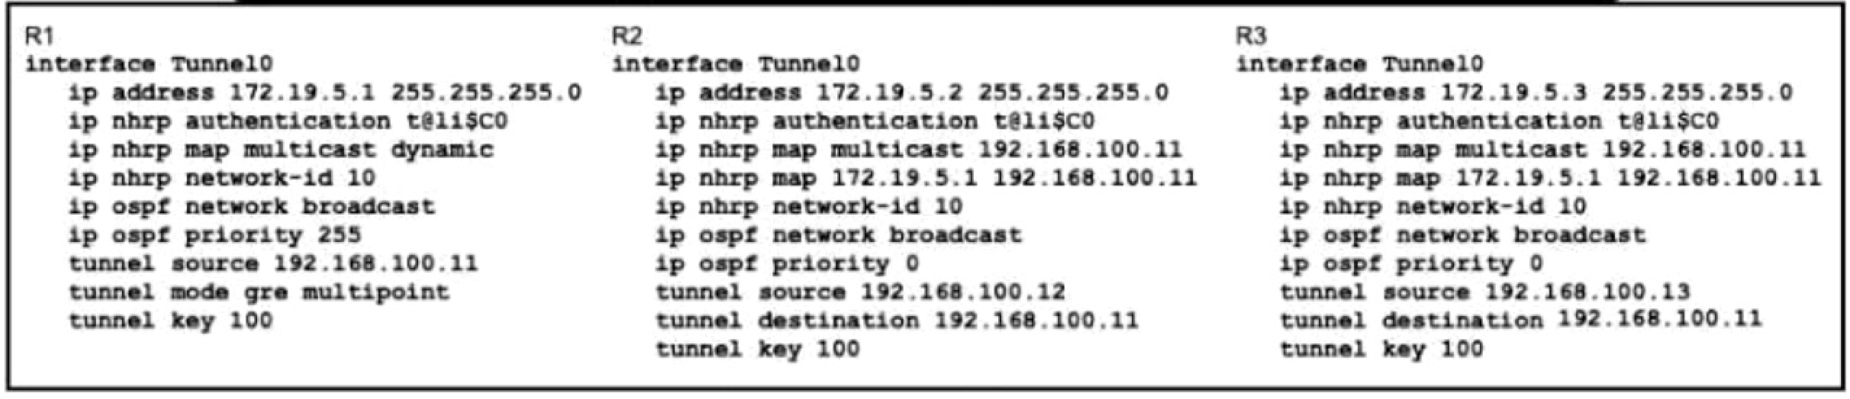

```
R1
interface Tunnel0
   ip address 172.19.5.1 255.255.255.0
   ip nhrp authentication t@li$C0
   ip nhrp map multicast dynamic
   ip nhrp network-id 10
   ip ospf network broadcast
   ip ospf priority 255
   tunnel source 192.168.100.11
   tunnel mode gre multipoint
   tunnel key 100
```

```
R2
interface Tunnel0
   ip address 172.19.5.2 255.255.255.0
   ip nhrp authentication t@li$C0
   ip nhrp map multicast 192.168.100.11
   ip nhrp map 172.19.5.1 192.168.100.11
   ip nhrp network-id 10
   ip ospf network broadcast
   ip ospf priority 0
   tunnel source 192.168.100.12
   tunnel destination 192.168.100.11
   tunnel key 100
```

```
R3
interface Tunnel0
   ip address 172.19.5.3 255.255.255.0
   ip nhrp authentication t@li$C0
   ip nhrp map multicast 192.168.100.11
   ip nhrp map 172.19.5.1 192.168.100.11
   ip nhrp network-id 10
   ip ospf network broadcast
   ip ospf priority 0
   tunnel source 192.168.100.13
   tunnel destination 192.168.100.11
   tunnel key 100
```

An organization is installing a new L3 MPLS link to establish DMVPN Phase 2 tunnels between the hub and two spoke routers. Which additional configuration should the engineer implement on each device to achieve optimal routing between the spokes?

A.
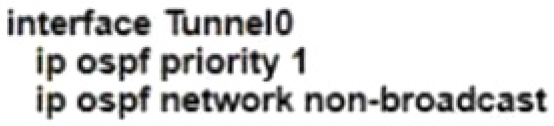

```
interface Tunnel0
  ip ospf priority 1
  ip ospf network non-broadcast
```

B.
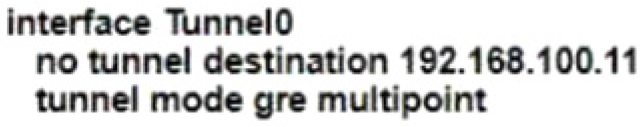

```
interface Tunnel0
  no tunnel destination 192.168.100.11
  tunnel mode gre multipoint
```

C.
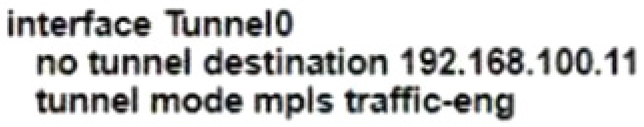

```
interface Tunnel0
  no tunnel destination 192.168.100.11
  tunnel mode mpls traffic-eng
```

D.
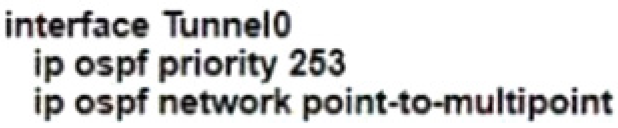

```
interface Tunnel0
  ip ospf priority 253
  ip ospf network point-to-multipoint
```

Correct Answer: *D*

⊟ **seal2** 3 months, 2 weeks ago

Selected Answer: B

agree, B is required for a good configuration

upvoted 1 times

⊟ **alex711** 3 months, 2 weeks ago

Selected Answer: B

I agree, B is correct.

upvoted 1 times

⊟ **Muste** 4 months, 1 week ago

Selected Answer: B

answer is B

upvoted 1 times

DenskyDen 4 months, 1 week ago

Selected Answer: B

Answer is B.

The question says its trying to establish a DMVPN Phase 2. This means we need to remove the static configuration of tunnel destination from the spokes and configure multipoint GRE.

upvoted 3 times

jansan55 4 months, 1 week ago

Selected Answer: B

My choice: answer B
There is no need for tunnel destination and we need gre multipoint tunnel mode
The OSPF priority of hub must be higher than the priority of spokes.
A good source:
https://networklessons.com/cisco/ccie-routing-switching/dmvpn-phase-2-basic-configuration
https://networklessons.com/cisco/ccie-routing-switching/dmvpn-phase-2-ospf-routing

upvoted 3 times

inteldarvid 4 months, 2 weeks ago

Selected Answer: D

The give anwser is correct

upvoted 1 times

Question #479 *Topic 1*

Refer to the exhibit.

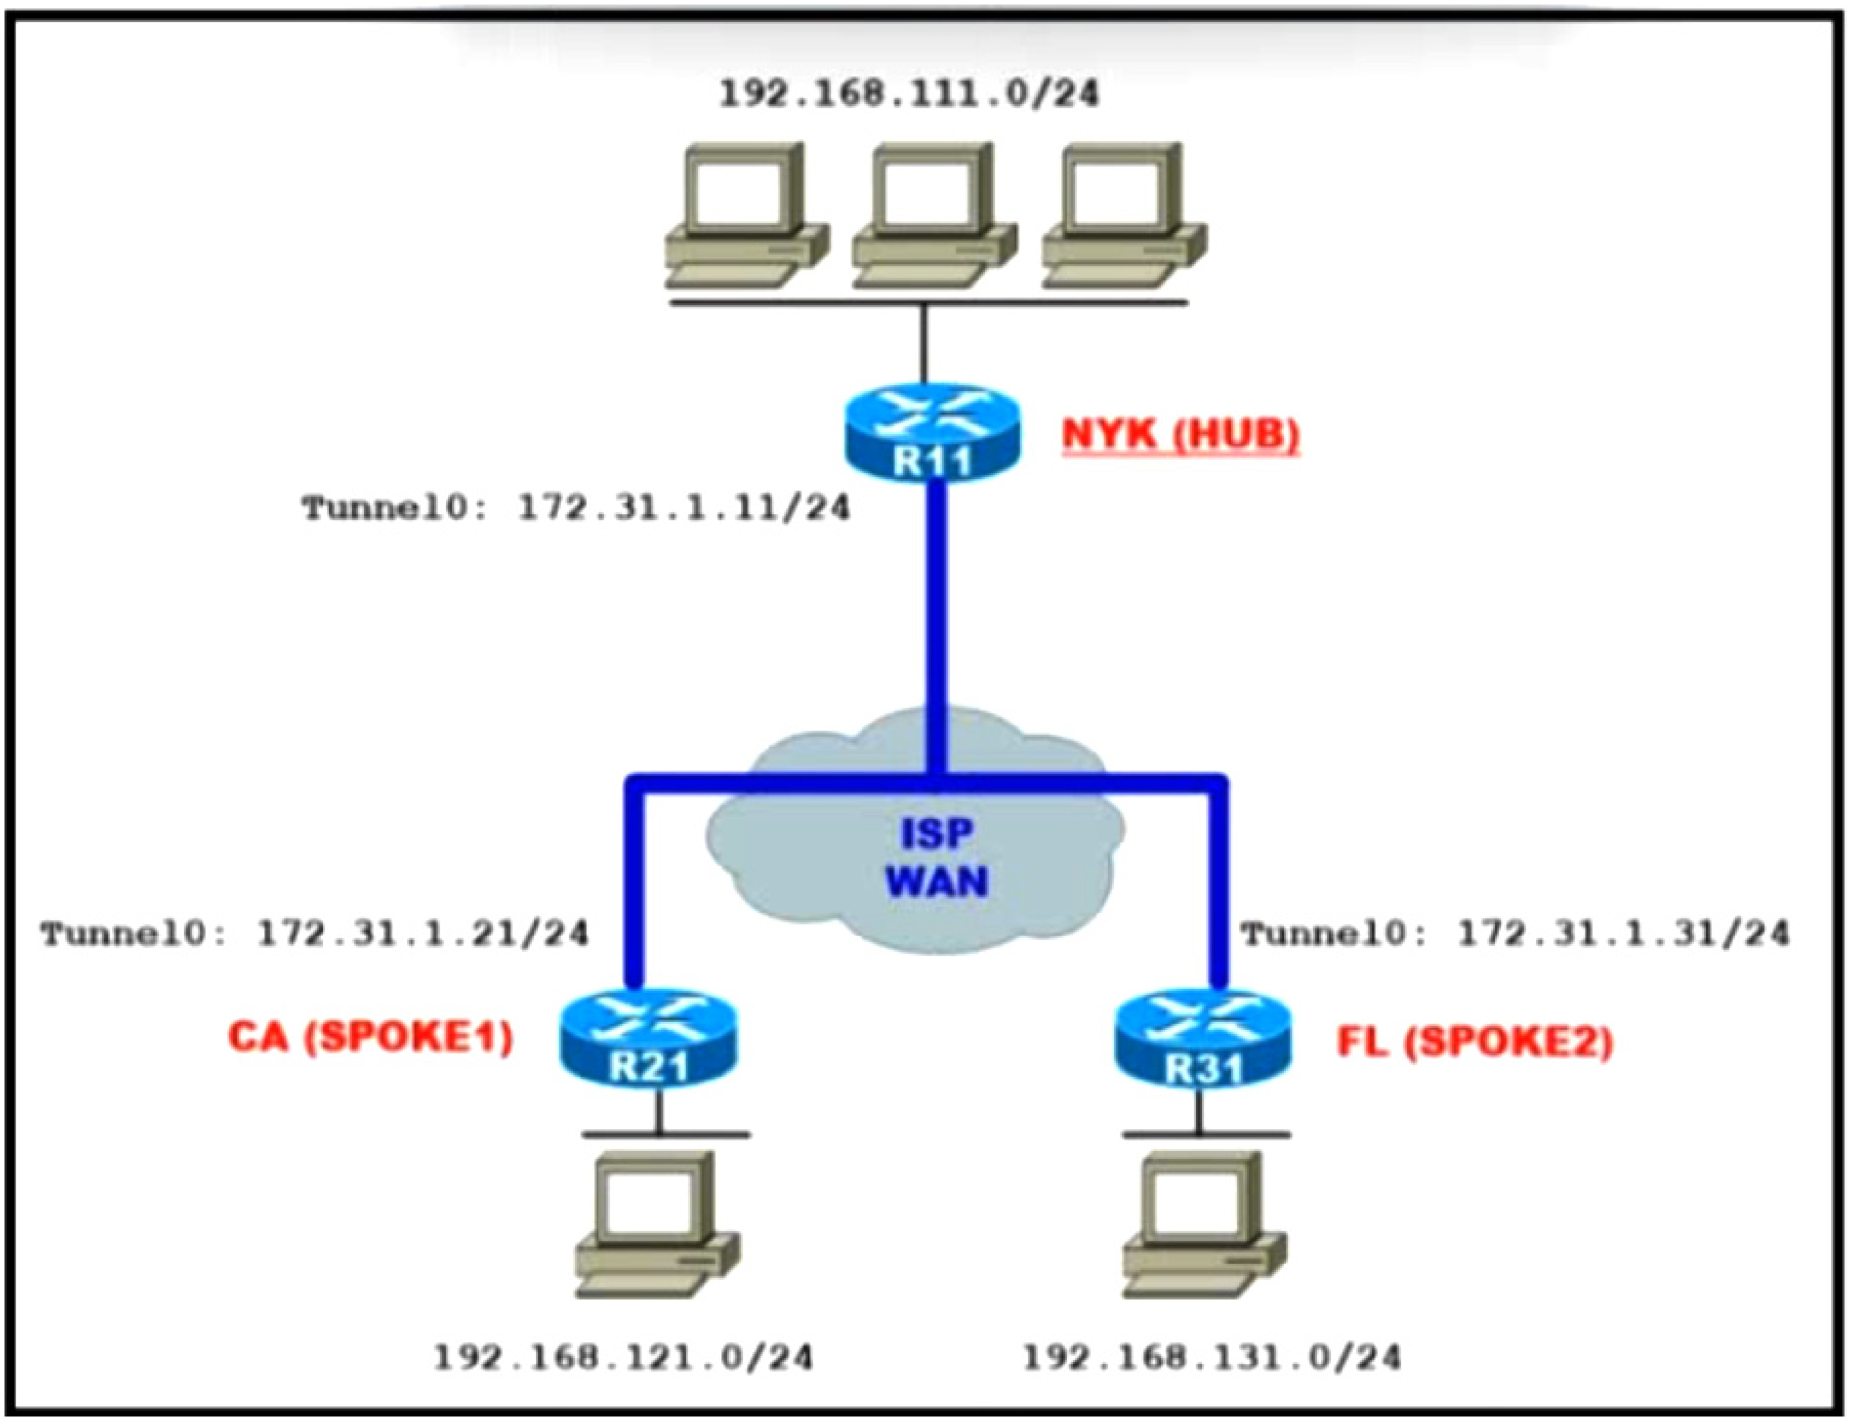

An engineer must configure the hub router to add new offices in the same infrastructure without performing any further configurations at the hub router.

Which tunnel mode configuration on the hub router meets this requirement?

A.
```
interface Tunnel0
 tunnel mode ipsec ipv4
```

B.
```
interface Tunnel0
 tunnel mode gre multipoint
```

C.
```
interface Tunnel0
 tunnel mode dvmrp
```

D.
```
interface Tunnel0
 tunnel mode ip
```

Correct Answer: *B*

seal2 3 months, 2 weeks ago

Selected Answer: B

Answer is correct, tunnel mode is gre multipoint for hub router

upvoted 1 times

wyattw 4 months, 2 weeks ago

Selected Answer: B

Answer is correct

upvoted 1 times

⊟ **inteldarvid** 4 months, 2 weeks ago

Selected Answer: B

Yes.correct

upvoted 1 times

SIMULATION

-

Guidelines

-

This is a lab item in which tasks will be performed on virtual devices.

• Refer to the Tasks tab to view the tasks for this lab item.
• Refer to the Topology tab to access the device console(s) and perform the tasks.
• Console access is available for all required devices by clicking the device icon or using the tab(s) above the console window.
• All necessary preconfigurations have been applied.
• Do not change the enable password or hostname for any device.
• Save your configurations to NVRAM before moving to the next item.
• Click Next at the bottom of the screen to submit this lab and move to the next question.
• When Next is clicked, the lab closes and cannot be reopened.

Topology

-

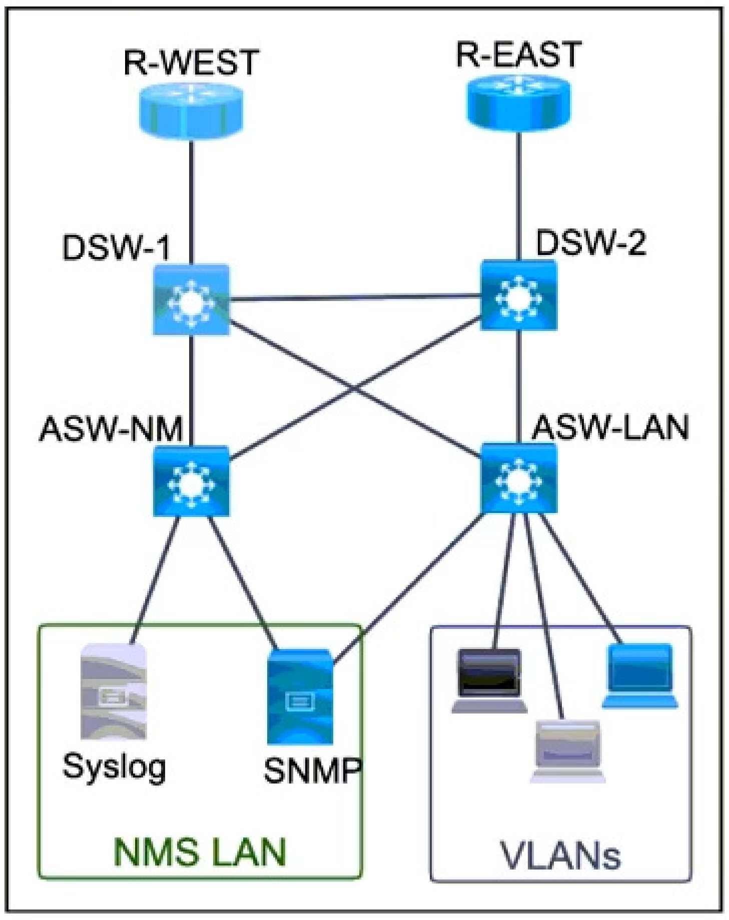

Tasks

-

Troubleshoot R-WEST to achieve the desired results:

1. The locally generated logs should have sequence numbers, date and time.

2. The SNMP traps related to OSPF and participating interface state changes utilizing RFC1253-MIB OSPFv2 should be sent to SNMP server.

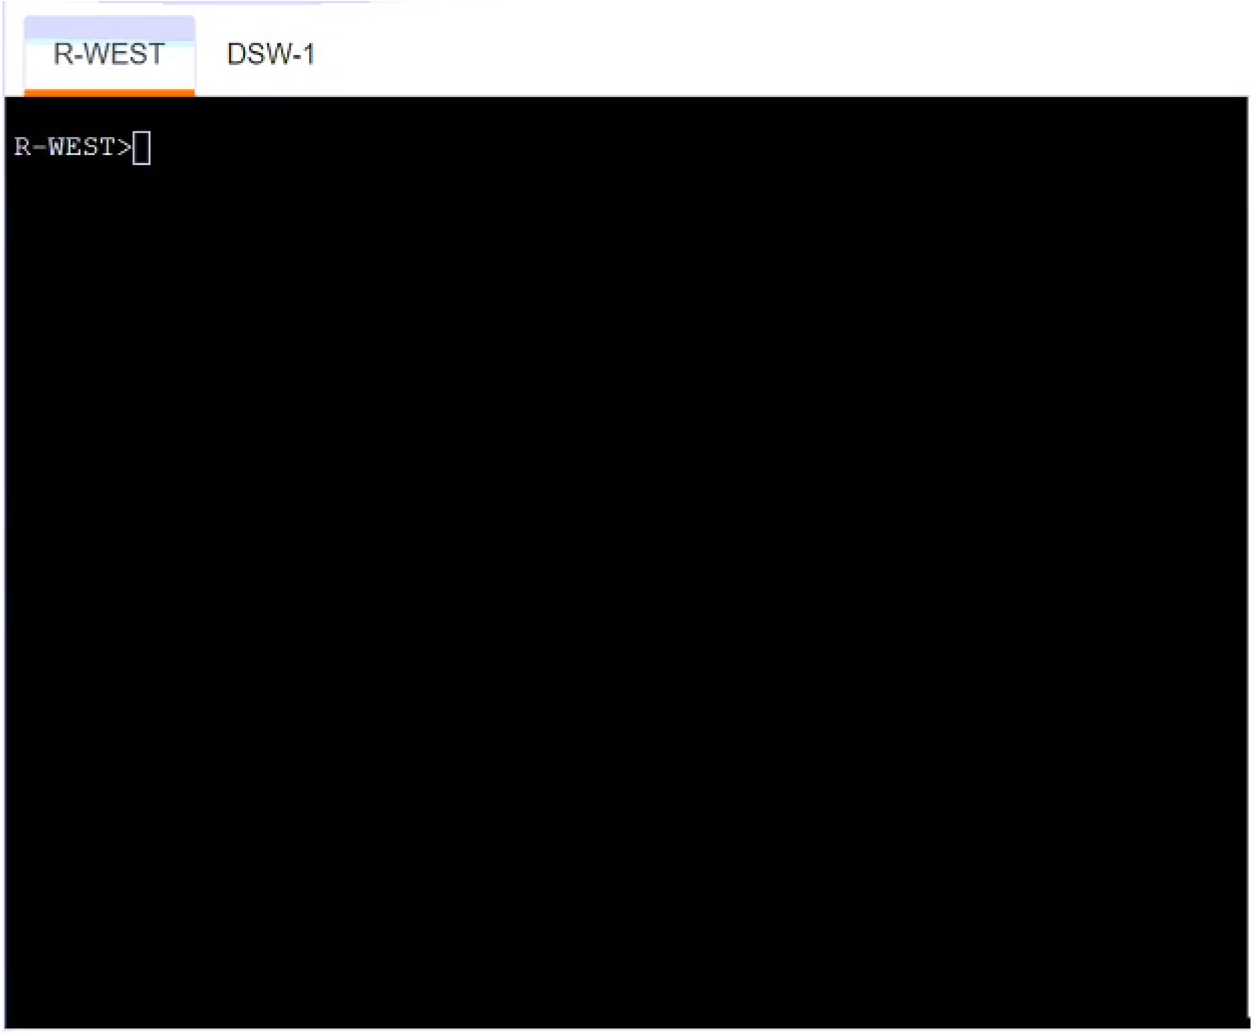

```
R-WEST>
```

Correct Answer:

```
R-WEST:
R-WEST>en
R-WEST#config t
R-WEST(config)#service sequence-numbers
R-WEST(config)#service timestamps log datetime msec
R-WEST(config)#snmp-server enable traps ospf
R-WEST(config)#end
R-WEST#write mem
```

mouin 3 months ago

R1(config)#service sequence-numbers

R1(config)#service timestamps log datetime

R1(config)#snmp-server enable traps ospf

upvoted 4 times

Question #481 *Topic 1*

Which two NLRI attributes are used by an MPLS Layer 3 VPN network to exchange VPNv4 routes between MPLS routers via MP-BGP? (Choose two.)

A. VPNv4 Prefix

B. Next Hop

C. Extended-Community

D. IPv4 Prefix

E. RT

Correct Answer: *BD*

⊟ **Brand** 3 months, 2 weeks ago

Selected Answer: BD

Provided answers are correct.

https://networklessons.com/mpls/mpls-layer-3-vpn-explained

upvoted 1 times

Question #482 Topic 1

Refer to the exhibit.

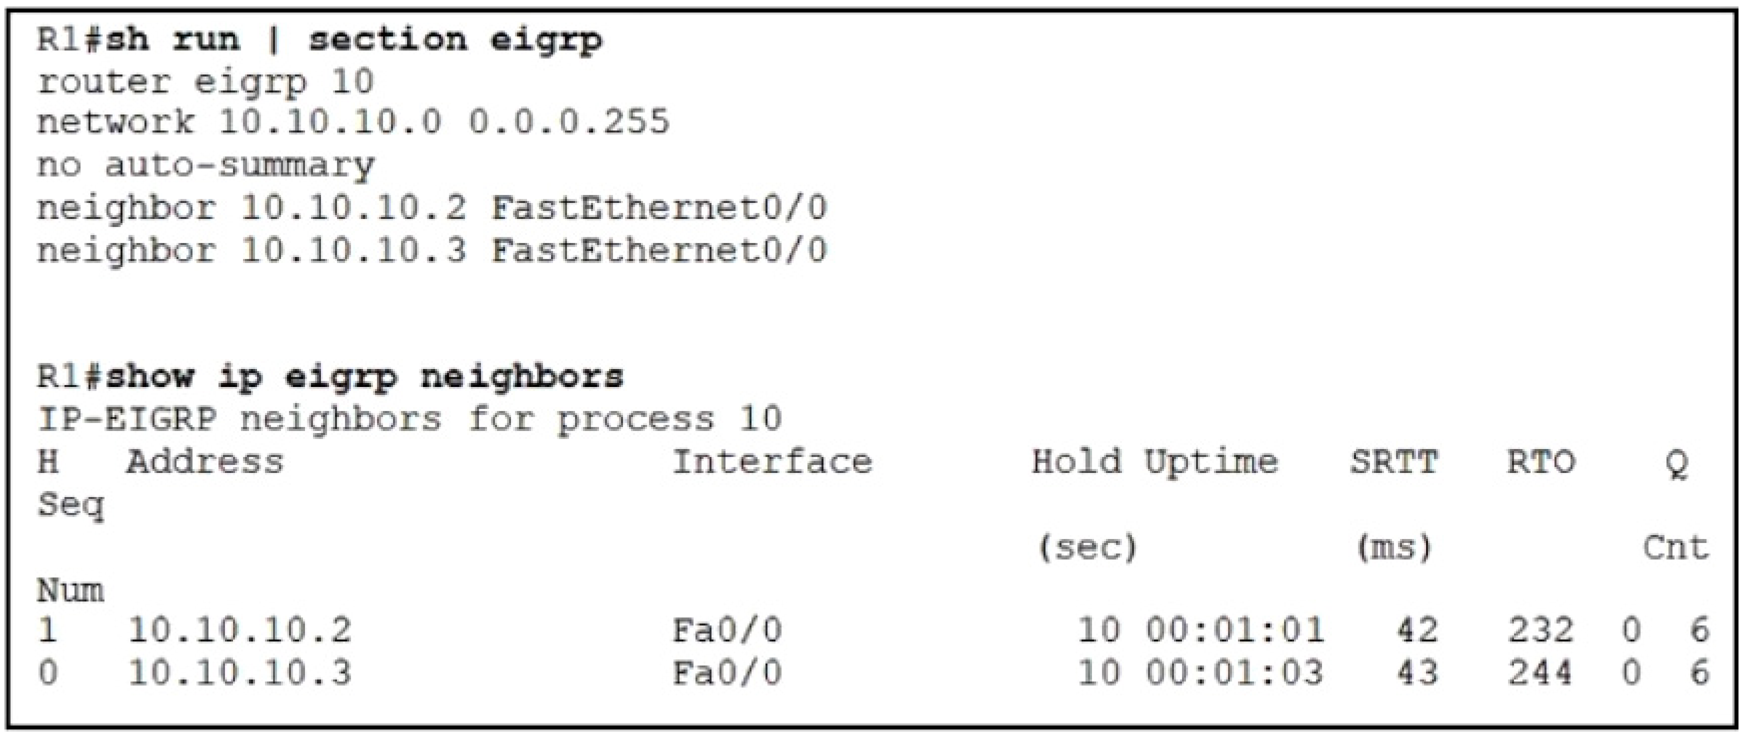

```
R1#sh run | section eigrp
router eigrp 10
network 10.10.10.0 0.0.0.255
no auto-summary
neighbor 10.10.10.2 FastEthernet0/0
neighbor 10.10.10.3 FastEthernet0/0


R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 10
H   Address                Interface        Hold Uptime   SRTT   RTO    Q
Seq
                                            (sec)         (ms)          Cnt
Num
1   10.10.10.2             Fa0/0              10 00:01:01   42   232   0   6
0   10.10.10.3             Fa0/0              10 00:01:03   43   244   0   6
```

The remote branch locations have a static neighbor relationship configured to R1 only. R1 has successful neighbor relationships with the remote locations of R2 and R3, but the end users cannot communicate with each other.

Which configuration resolves the issue?

A.

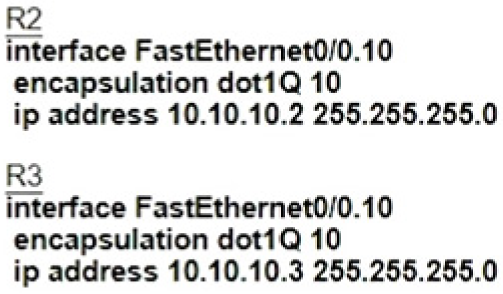

```
R2
interface FastEthernet0/0.10
 encapsulation dot1Q 10
 ip address 10.10.10.2 255.255.255.0
R3
interface FastEthernet0/0.10
 encapsulation dot1Q 10
 ip address 10.10.10.3 255.255.255.0
```

B.

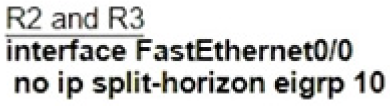

```
R2 and R3
interface FastEthernet0/0
 no ip split-horizon eigrp 10
```

C.

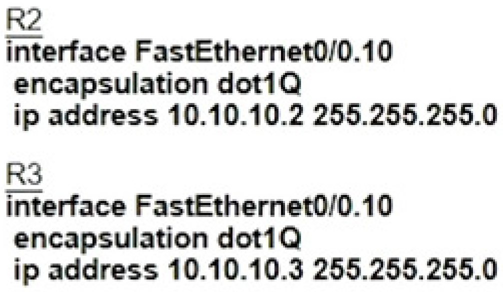

```
R2
interface FastEthernet0/0.10
 encapsulation dot1Q
 ip address 10.10.10.2 255.255.255.0
R3
interface FastEthernet0/0.10
 encapsulation dot1Q
 ip address 10.10.10.3 255.255.255.0
```

D.

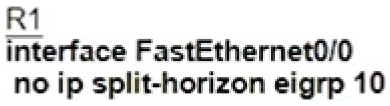

```
R1
interface FastEthernet0/0
 no ip split-horizon eigrp 10
```

Correct Answer: *D*

seal2 3 months, 2 weeks ago

Selected Answer: D

split-horizon is default behavior for EIGRP, so it will not show up in the configuration. it stops the spokes from having routes to each other, so we need to disable it.

upvoted 2 times

Question #483 *Topic 1*

Refer to the exhibit.

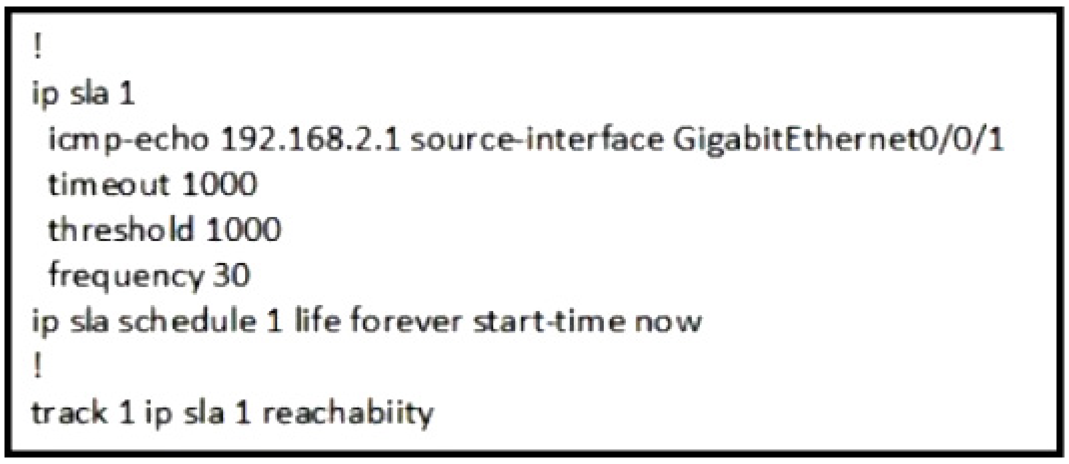

```
!
ip sla 1
 icmp-echo 192.168.2.1 source-interface GigabitEthernet0/0/1
 timeout 1000
 threshold 1000
 frequency 30
ip sla schedule 1 life forever start-time now
!
track 1 ip sla 1 reachabiity
```

An engineer observes that every time the ICMP packet is lost at a polling interval, track 1 goes down, which causes unnecessary disruption and instability in the network. The engineer does not want the traffic to be rerouted if the loss of ICMP packets is negligible. If the packet loss is persistent for a longer duration, the track must go down and the traffic must be rerouted. Which action resolves the issue?

A. Change the IP SLA schedule to run only at certain intervals.

B. Increase the timeout value from 1000 to 1500.

C. Define a delay timer under track 1.

D. Increase the threshold value from 1000 to 1500.

Correct Answer: *C*

ZamanR 5 days, 5 hours ago

C is the answer

upvoted 1 times

RouterToRooter 2 weeks, 5 days ago

Selected Answer: C

Up/down Delay timer

upvoted 2 times

chris110 3 months, 1 week ago

Selected Answer: D

IP SLA parameters:

threshold – specifies after how much time a "reaction event" takes place, such as considering a particular SLA as "down." In simpler terms, for an icmp-echo SLA, it is the amount of time that must pass without a successful ping before the SLA is considered down.

timeout – the amount of time an IP SLA will wait for a response from its echo request packet.

-> Therefore we must increase the "threshold" parameter, not "timeout".

upvoted 1 times

seal2 3 months, 2 weeks ago

Selected Answer: C

Correct, a delay timer configured under the IP SLA will give leniency so that track 1 won't go down if packet loss is small enough.

upvoted 1 times

robi1020 3 months, 2 weeks ago

Selected Answer: C

Correct!

https://community.cisco.com/t5/switching/delay-down-timer-frequency-timer-interaction-with-ip-sla-during/td-p/2808130

upvoted 1 times

Refer to the exhibit.

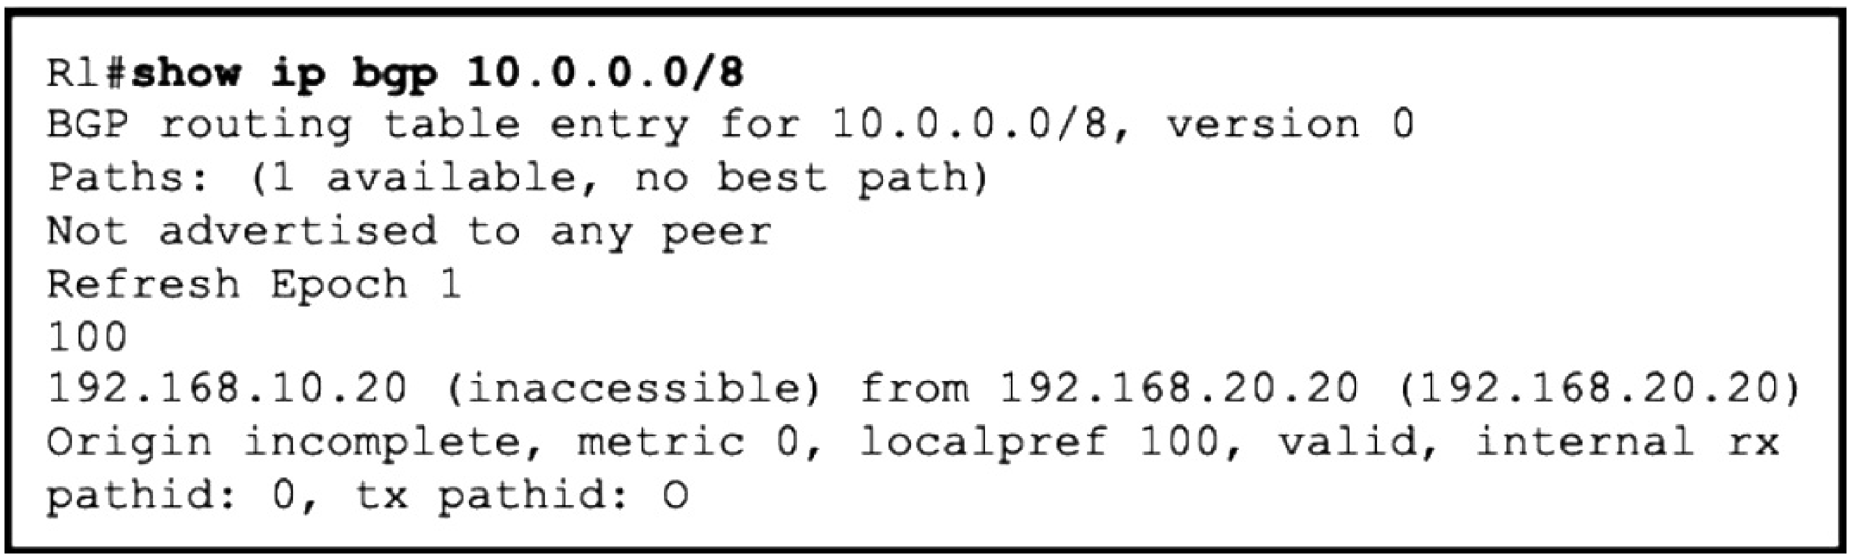

```
R1#show ip bgp 10.0.0.0/8
BGP routing table entry for 10.0.0.0/8, version 0
Paths: (1 available, no best path)
Not advertised to any peer
Refresh Epoch 1
100
192.168.10.20 (inaccessible) from 192.168.20.20 (192.168.20.20)
Origin incomplete, metric 0, localpref 100, valid, internal rx
pathid: 0, tx pathid: 0
```

An engineer is troubleshooting a prefix advertisement issue from R3, which is not directly connected to R1. Which configuration resolves the issue?

A. R2(config)#router bgp 64512 -
R2(config-router)#neighbor 192.168.20.10 next-hop-self

B. R1(config)#router bgp 64512 -
R1(config-router)#neighbor 192.168.10.20 next-hop-self

C. R1(config)#router bgp 64512 -
R1(config-router)#neighbor 192.168.20.20 next-hop-self

D. R2(config)#router bgp 64512 -
R2(config-router)#neighbor 192.168.10.20 next-hop-self

Correct Answer: *C*

Storcaks Highly Voted 3 months, 1 week ago

Selected Answer: A

C doesn't really make any sense.
Next-hop-self needs to be enabled on router 192.168.20.20 which is presumably R2, not on R1 itself. R2 (192.168.20.20) is peering with R3 (192.168.10.20) and R1 (192.168.20.10).
I'd go with A

upvoted 5 times

RouterToRooter Most Recent 1 week, 5 days ago

Selected Answer: C

Correct answer is C for me

upvoted 1 times

sayed_2908 3 weeks, 5 days ago

answer is C

R3 <------> R2 <------> R1. R1 is not directly connected to R3

192.168.10.20 is R3
192.168.20.20 is R2
192.168.20.10 is R1

next-hop self need to be set on R2

upvoted 1 times

chris110 3 months, 1 week ago

Take a look Q95

https://www.networktut.com/new-enarsi-questions-6

The next hop is preserved in eBGP advertisements for IBGP neighbor so it will not change when R2 advertises to IBGP router R1 and this is the case of this question. The below figure helps you understand it:

BGP_inaccessible_next_hop_self.jpg

Although R2 keeps the next-hop 192.168.10.20 when advertising to R1 for network 10.0.0.0/8 but R1 does not know how to reach this next-hop so it is marked as inaccessible.

Also from the output "192.168.10.20 (inaccessible) from 192.168.20.20", we learn that the IP address 192.168.10.20 was learned from 192.168.20.20 so this IP address belongs to R2 interface which is facing R1. Therefore the IP address of R1 interface which is facing R2 should be 192.168.20.10.

One solution of this problem is to set the command "neighbor 192.168.20.10 next-hop-self" on R2 so that R2 advertise its IP address as the next-hop for R1.

upvoted 1 times

Question #485 Topic 1

Refer to the exhibit.

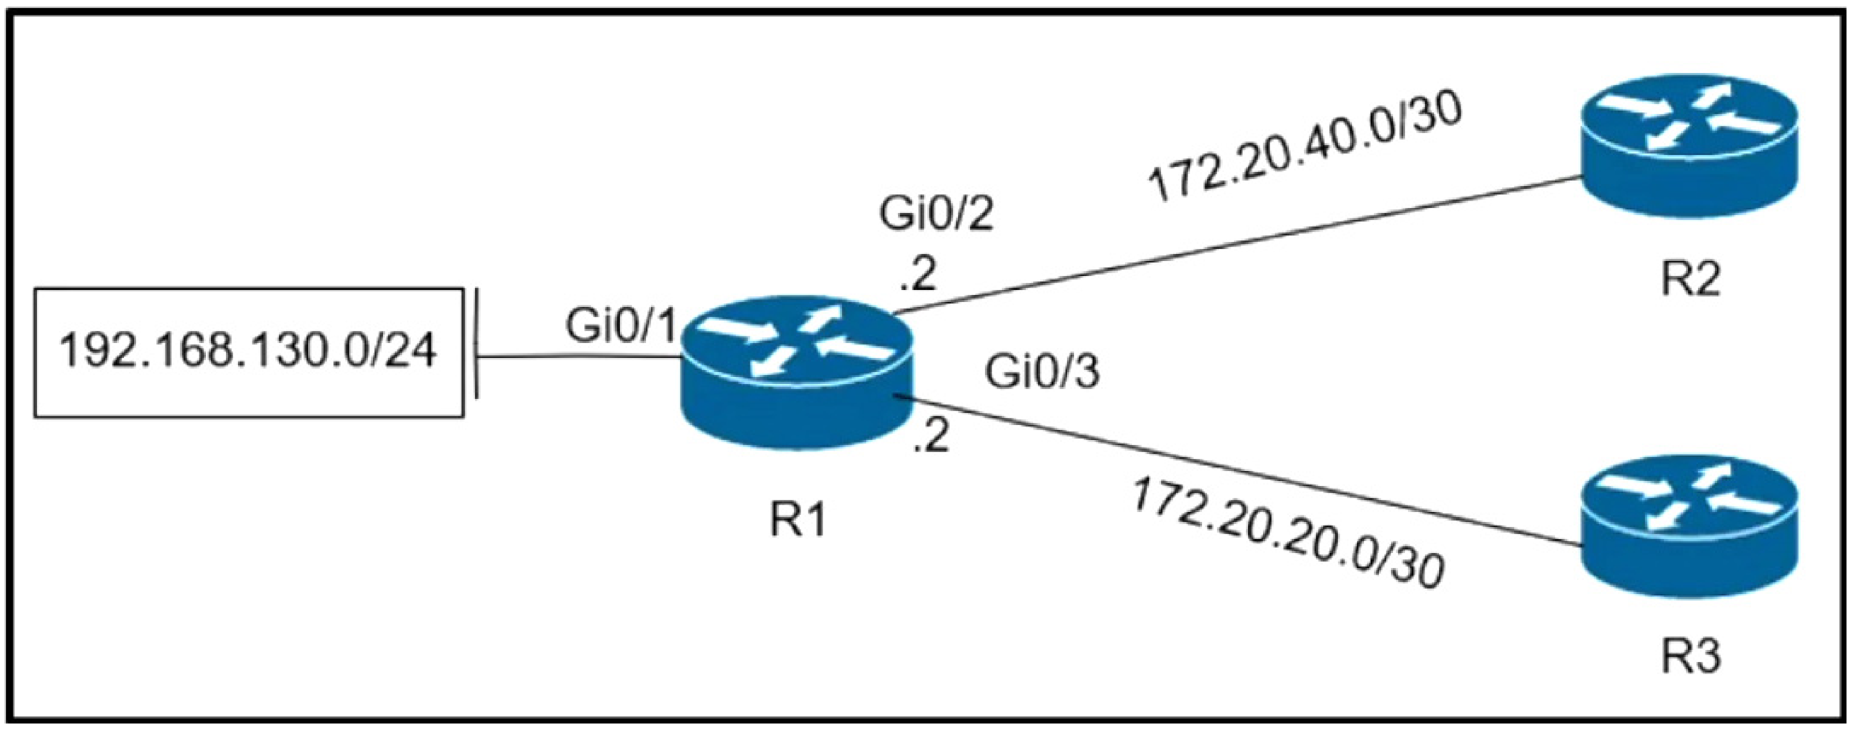

Which policy configuration on R1 forwards any traffic that is sourced from the 192.168.130.0/24 network to R2?

A.

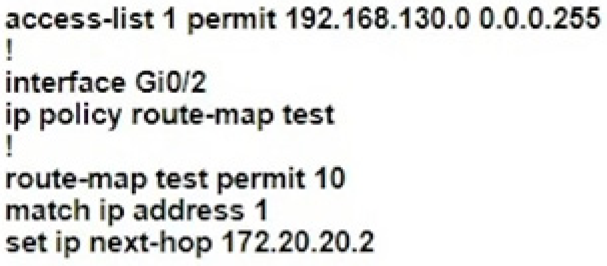

```
access-list 1 permit 192.168.130.0 0.0.0.255
!
interface Gi0/2
ip policy route-map test
!
route-map test permit 10
match ip address 1
set ip next-hop 172.20.20.2
```

B.

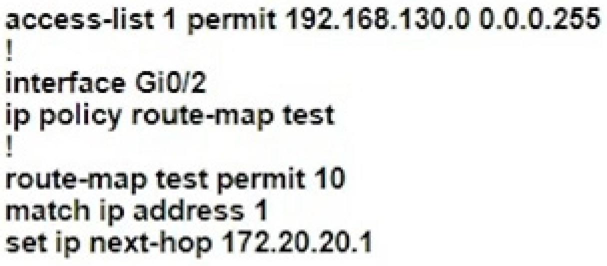

```
access-list 1 permit 192.168.130.0 0.0.0.255
!
interface Gi0/2
ip policy route-map test
!
route-map test permit 10
match ip address 1
set ip next-hop 172.20.20.1
```

C.

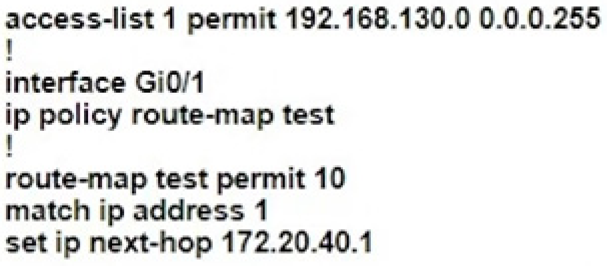

```
access-list 1 permit 192.168.130.0 0.0.0.255
!
interface Gi0/1
ip policy route-map test
!
route-map test permit 10
match ip address 1
set ip next-hop 172.20.40.1
```

D.

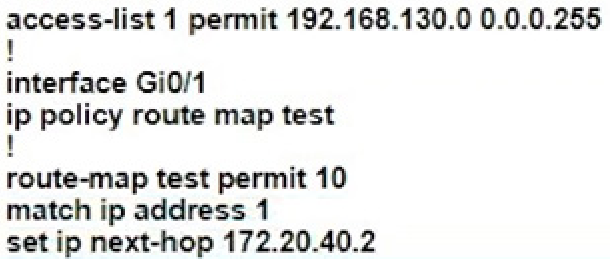

```
access-list 1 permit 192.168.130.0 0.0.0.255
!
interface Gi0/1
ip policy route map test
!
route-map test permit 10
match ip address 1
set ip next-hop 172.20.40.2
```

Correct Answer: *C*

seal2 3 months, 2 weeks ago

Selected Answer: C

C is correct, it's the only one that specifies the correct next hop.

upvoted 1 times

chris110 3 months, 2 weeks ago

Selected Answer: C

Its C, Next Hop ...40.1

upvoted 1 times

[Removed] 3 months, 2 weeks ago

Selected Answer: C

Correct, the ACL identifies the source network correctly. The Route map references the ACL and identifies the next hop IPv4 correctly

upvoted 1 times

⊟ **Colmenarez** 3 months, 1 week ago

lol you still around here? I thought you had already taken the exam.

upvoted 1 times

Question #486 *Topic 1*

Refer to the exhibit.

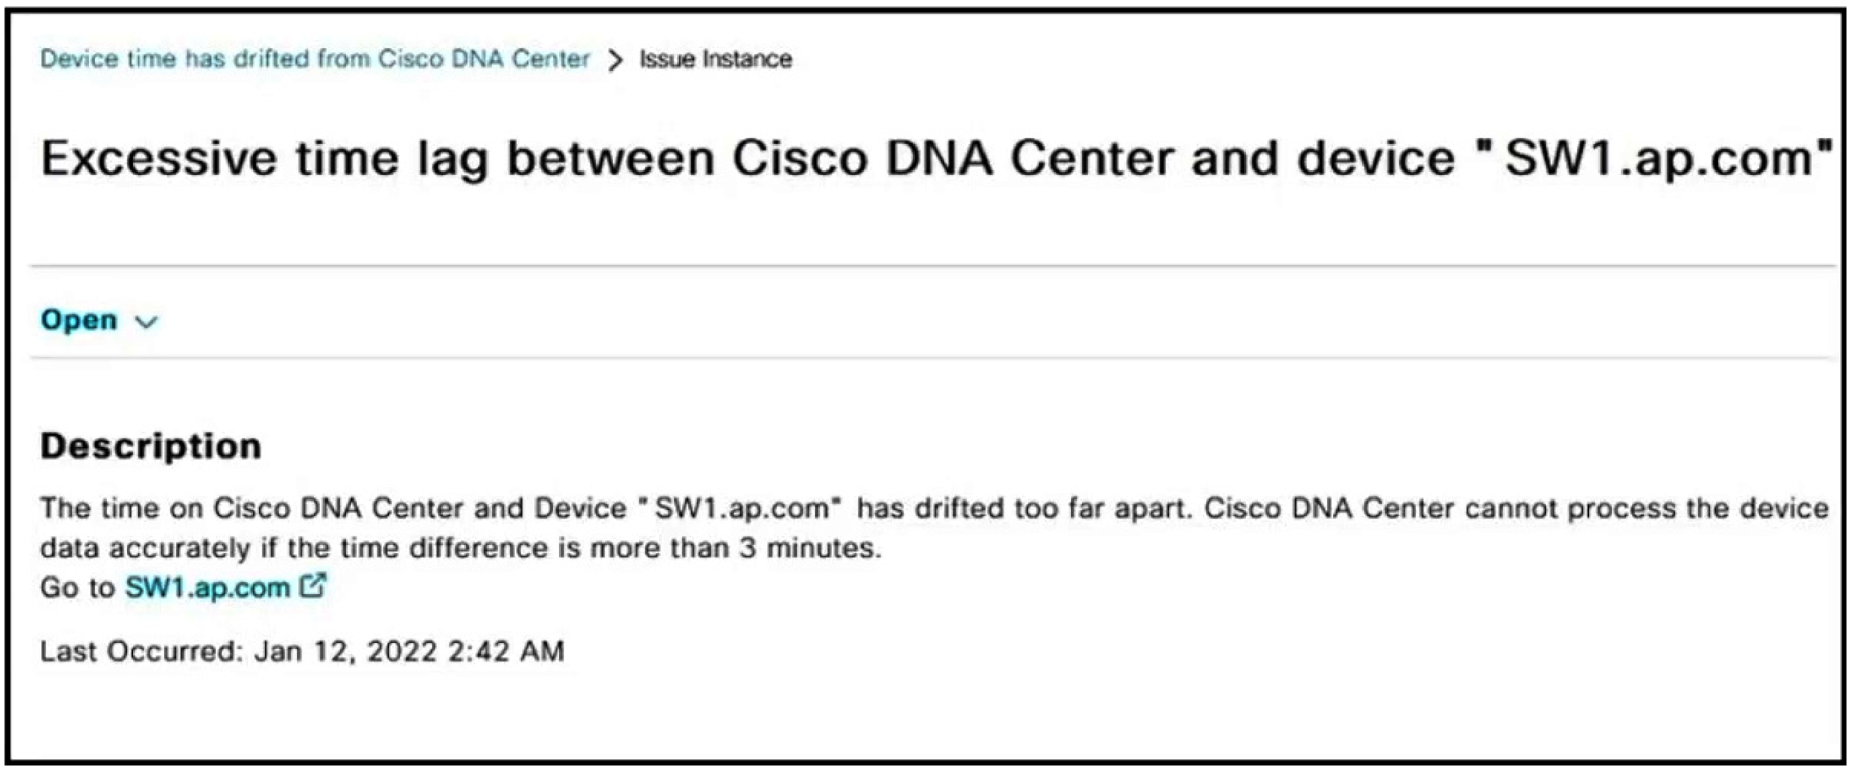

Device time has drifted from Cisco DNA Center > Issue Instance

Excessive time lag between Cisco DNA Center and device " SW1.ap.com"

Open

Description

The time on Cisco DNA Center and Device " SW1.ap.com" has drifted too far apart. Cisco DNA Center cannot process the device data accurately if the time difference is more than 3 minutes.
Go to SW1.ap.com

Last Occurred: Jan 12, 2022 2:42 AM

Which action resolves the issue?

A. Establish connectivity between the NTP server and the switch.

B. Configure the local time on the SW1 device.

C. Configure the local time on Cisco DNA Center.

D. Establish connectivity between the NTP server and Cisco DNA Center.

Correct Answer: *A*

⊟ **seal2** 3 months, 2 weeks ago

Selected Answer: A

A. an NTP server fixes this problem by not allowing the time to drift.
B. local time will drift eventually and the problem will resume
C. DNA center is unable to be modified in this way
D. same as C

upvoted 1 times

Question #487 *Topic 1*

What does IPv6 Source Guard utilize to determine if IPv6 source addresses should be forwarded?

A. ACLs

B. ACE

C. DHCP

D. Binding Table

Correct Answer: *D*

seal2 3 months, 2 weeks ago

A. source guard doesn't use ACLs
B. source guard doesn't use ACE
C. source guard does use DHCP information to fill a binding table, but is not directly responsible for source guard's actions
D. source guard uses the binding table for its duties. :)

upvoted 2 times

Question #488 *Topic 1*

Refer to the exhibit.

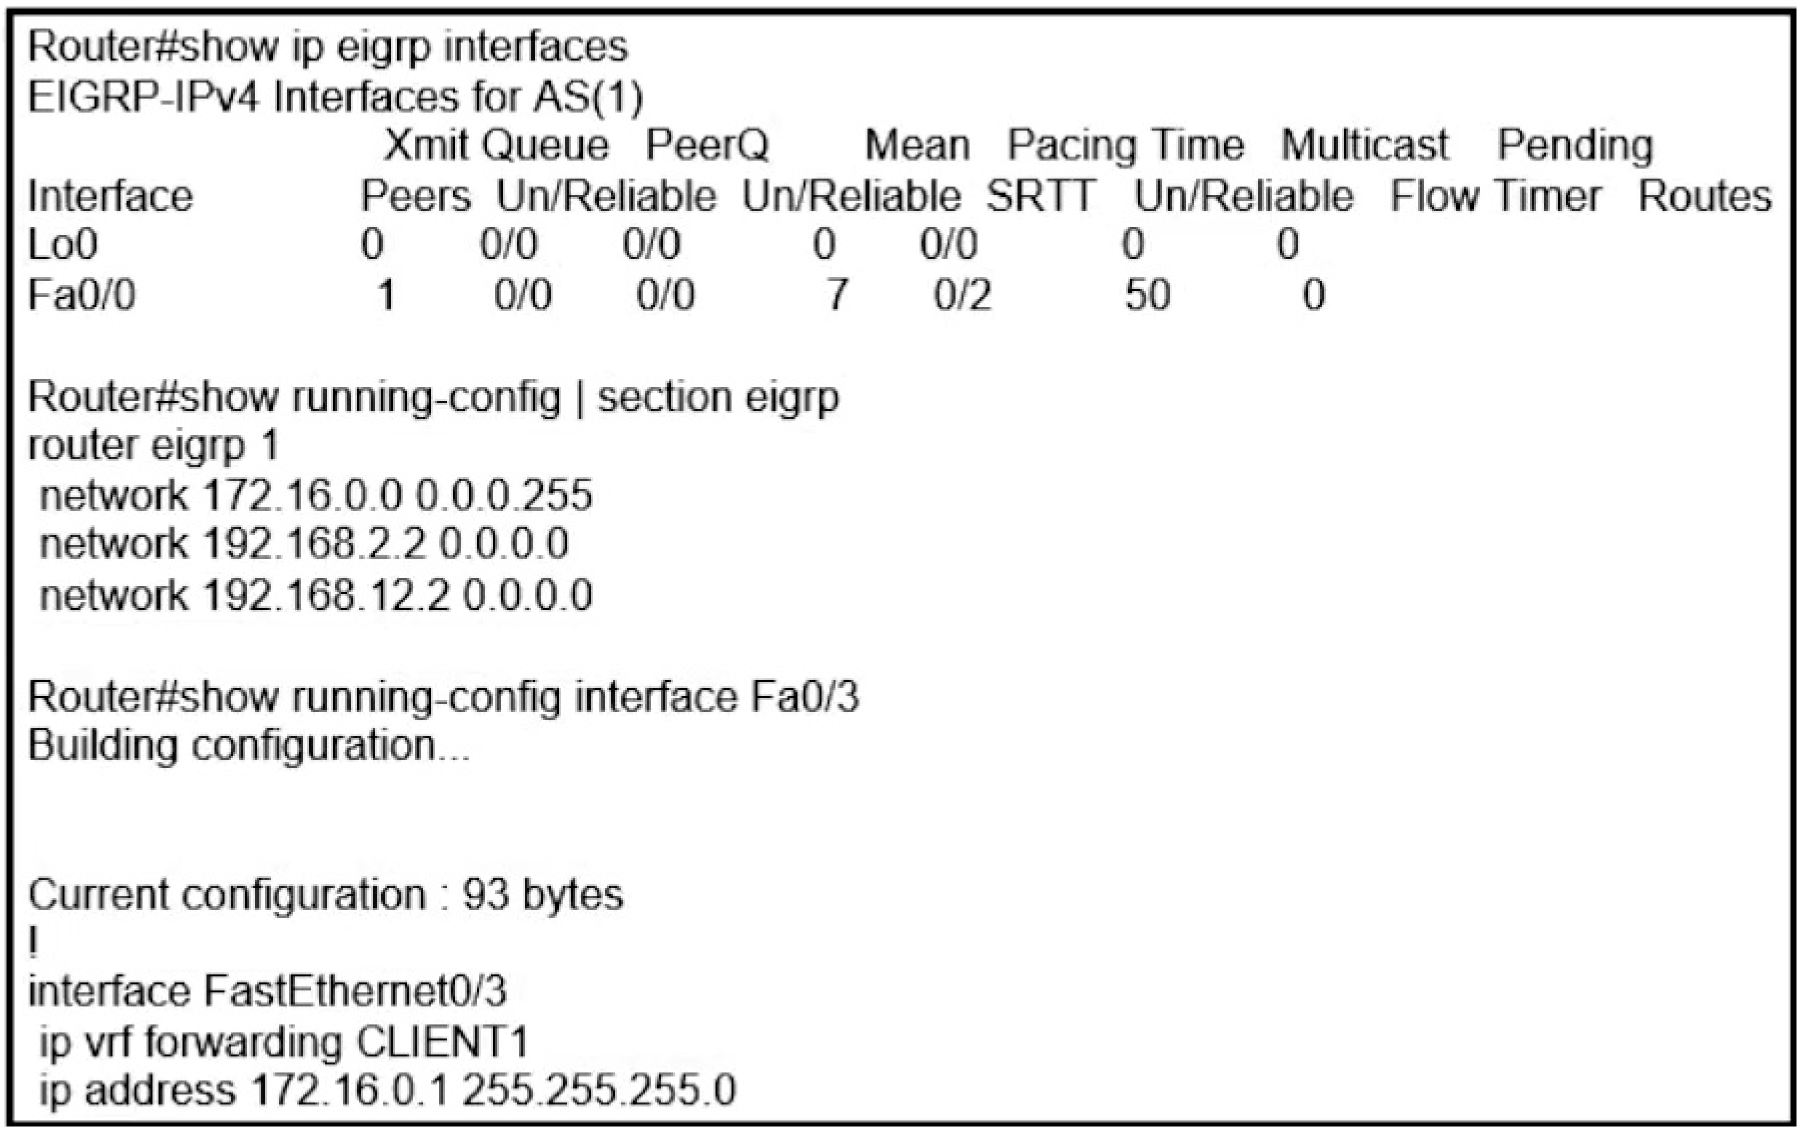

```
Router#show ip eigrp interfaces
EIGRP-IPv4 Interfaces for AS(1)
                     Xmit Queue   PeerQ       Mean  Pacing Time  Multicast    Pending
Interface           Peers  Un/Reliable  Un/Reliable  SRTT  Un/Reliable  Flow Timer  Routes
Lo0                 0       0/0       0/0         0      0/0          0          0
Fa0/0                1       0/0       0/0         7      0/2          50         0

Router#show running-config | section eigrp
router eigrp 1
 network 172.16.0.0 0.0.0.255
 network 192.168.2.2 0.0.0.0
 network 192.168.12.2 0.0.0.0

Router#show running-config interface Fa0/3
Building configuration...


Current configuration : 93 bytes
!
interface FastEthernet0/3
 ip vrf forwarding CLIENT1
 ip address 172.16.0.1 255.255.255.0
```

While troubleshooting an EIGRP neighbor adjacency problem, the network engineer notices that the interface connected to the neighboring router is not participating in the EIGRP process. Which action resolves the issue?

A. Configure EIGRP metrics on interface FastEthemet0/3.

B. Configure the network command under EIGRP address family vrf CLIENT1.

C. Configure the network command under EIGRP address family ipv4.

D. Configure the network command to network 172.16.0.1 0.0.0.0.

Correct Answer: *B*

ZamanR 1 week, 4 days ago

B Answer is correct
router eigrp 1
!

address-family ipv4 vrf CLIENT1

network 172.16.0.0 0.0.0.255

no auto-summary

autonomous-system 1

exit-address-family

upvoted 1 times

DeWalt95 2 weeks, 4 days ago

Actual answer is a combination of B and C.

Address-family ipv4 vrf 'name'

upvoted 1 times

Tester948 4 weeks ago

Selected Answer: C

Should be C:

```
R7(config)#router eigrp 100
R7(config-router)#add
R7(config-router)#address-family ?
ipv4 Address family

R7(config-router)#address-family ip
R7(config-router)#address-family ipv4 ?
unicast Address Family Modifier
vrf Specify parameters for a VPN Routing/Forwarding instance
<cr>

R7(config-router)#address-family ipv4 vr
R7(config-router)#address-family ipv4 vrf ?
WORD VPN Routing/Forwarding instance name
```

upvoted 1 times

Question #489 *Topic 1*

What are two characteristics of a VRF instance? (Choose two.)

A. A customer site can be associated to different VRFs.

B. Each VRF has a different set of routing and CEF tables.

C. It is defined by the VPN membership of a customer site attached to a P device.

D. All VRFs share customers routing and CEF tables.

E. An interface must be associated to one VRF.

Correct Answer: *BE*

seal2 3 months, 2 weeks ago

Selected Answer: BE

A. no, one VRF per site should suffice
B. this is kind of the point of VRFs. true
C. no, customer sites are not directly connected to P devices, so i don't consider this true
D. no, this is what we avoid by using vrfs
E. yes, you cannot associate multiple VRFs to one interface.

upvoted 1 times

chris110 3 months, 2 weeks ago

Selected Answer: BE

Two characteristics of a VRF (Virtual Routing and Forwarding) instance are:

B. Each VRF has a different set of routing and CEF tables: VRFs maintain separate routing and CEF (Cisco Express Forwarding) tables, allowing different VRFs to have isolated routing information and forwarding decisions.

E. An interface must be associated with one VRF: Interfaces on a router or switch are typically associated with a specific VRF. This association ensures that traffic on that interface is segregated and follows the routing information within the designated VRF.

So, the correct options are B and E.

upvoted 1 times

Question #490

Topic 1

How is the LDP router ID used in an MPLS network?

A. The force keyword changes the router ID to the specified address without causing any impact.

B. The loopback with the highest IP address is selected as the router ID.

C. The MPLS LDP router ID must match the IGP router ID.

D. If not configured, the operational physical interface is chosen as the router ID even if a loopback is configured.

Correct Answer: *B*

seal2 3 months, 2 weeks ago

Selected Answer: B

A. not true, force causes downtime
B. true, the highest loopback is chosen if not manually configured, then if there are no loopbacks it chooses the highest IP address given to any other interface.
C. doesn't really make sense to necessitate them matching anyways, it's not true.
D. no, it selects highest loopback first, then if unsuccessful it tries physical interfaces.

upvoted 1 times

chris110 3 months, 2 weeks ago

Selected Answer: B

In an MPLS (Multiprotocol Label Switching) network, the LDP (Label Distribution Protocol) router ID is used to uniquely identify a router within the MPLS domain. The correct statement about the LDP router ID is:

B. The loopback with the highest IP address is selected as the router ID: By default, the LDP router ID is determined based on the loopback interface with the highest IP address configured on the router. This provides a stable and predictable router ID that is not tied to the operational state of physical interfaces.

The other options (A, C, and D) do not accurately describe the default behavior of how the LDP router ID is determined in an MPLS network.

upvoted 1 times

Refer to the exhibit.

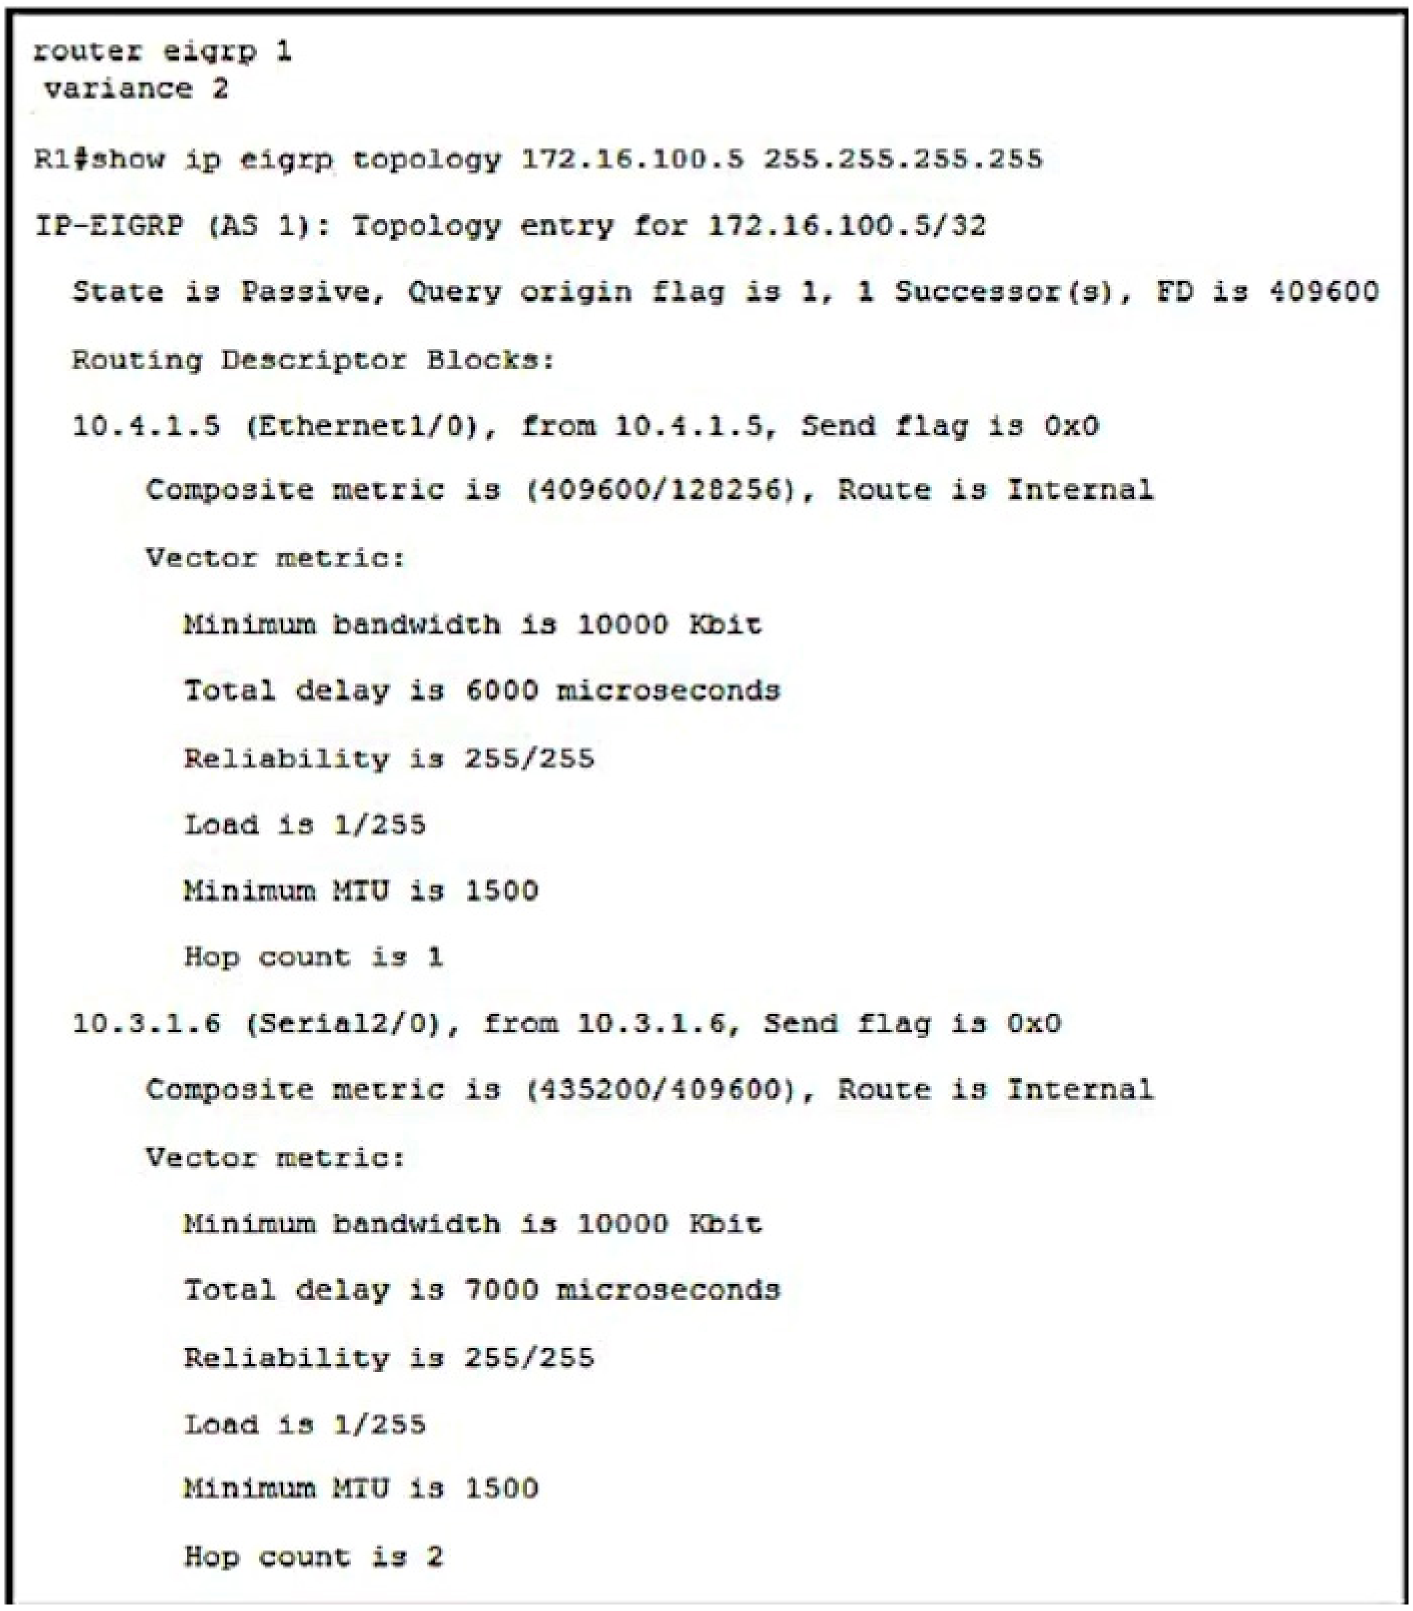

```
router eigrp 1
 variance 2

R1#show ip eigrp topology 172.16.100.5 255.255.255.255

IP-EIGRP (AS 1): Topology entry for 172.16.100.5/32

   State is Passive, Query origin flag is 1, 1 Successor(s), FD is 409600

   Routing Descriptor Blocks:

   10.4.1.5 (Ethernet1/0), from 10.4.1.5, Send flag is 0x0

       Composite metric is (409600/128256), Route is Internal

       Vector metric:

         Minimum bandwidth is 10000 Kbit

         Total delay is 6000 microseconds

         Reliability is 255/255

         Load is 1/255

         Minimum MTU is 1500

         Hop count is 1

   10.3.1.6 (Serial2/0), from 10.3.1.6, Send flag is 0x0

       Composite metric is (435200/409600), Route is Internal

       Vector metric:

         Minimum bandwidth is 10000 Kbit

         Total delay is 7000 microseconds

         Reliability is 255/255

         Load is 1/255

         Minimum MTU is 1500

         Hop count is 2
```

A network engineer troubleshooting a packet drop problem for the host 172.16.100.5 notices that only one link is used and installed on the routing table, which saturates the bandwidth. Which action must the engineer take to resolve the high bandwidth utilization problem and share the traffic toward this host between the two available links?

A. Disable the eigrp split horizon loop protection mechanism.

B. Set the eigrp variance equal to 3 to install a second route with a metric not larger than 3 times of the best metric.

C. Change the EIGRP delay metric to meet the feasibility condition.

D. Set the eigrp variance equal to 4 to install a second route with a metric not larger than 4 times of the best metric.

Correct Answer: *D*

AlexInShort12 4 days, 7 hours ago

Selected Answer: B

With the variance to 2, the RT should already load balance between the two routes...
The feasibility condition doesn't talk about the delay only the FD-AD....
Probably something wrong with the question.

upvoted 1 times

louisvuitton12 1 month, 2 weeks ago

Selected Answer: C

It is C for me

upvoted 1 times

⊟ **chris110** 3 months, 1 week ago

Selected Answer: C

From the output of the "show ip eigrp topology ..." command, we notice network 172.16.100.5/32 was learned via two routes:
+ From 10.4.1.5 with FD = 409600 and AD = 128256
+ From 10.3.1.6 with FD = 435200 and AD = 409600
To use both paths (called unequal cost load balancing) with EIGRP, the second path must satisfy the feasibility condition. The feasibility condition states that, the Advertised Distance (AD) of a route must be lower than the feasible distance of the current successor route.
In this case, the second path did not satisfy the feasible condition as its AD (409600) is equal to the FD of the best path -> Therefore we cannot configure load balancing with "variance" command.

upvoted 1 times

⊟ **chris110** 3 months, 1 week ago

The only reasonable solution of this question is "change the delay metric" so that the value of the FD of the best path is higher than its current value to meet the feasibility condition.
Suppose after changing the delay metric, the second path now met feasibility condition. Let's check if the second path would be installed into the routing table:
The EIGRP will install all paths with metric < variance * best_metric into the local routing table. Therefore we can calculate the variance > metric / best_metric = 435200 / 409600 =1.06 -> A variance of 2 is enough to make EIGRP install the second path to its routing table.

upvoted 1 times

⊟ **Colmenarez** 3 months, 2 weeks ago

Selected Answer: C

variance 2 is already in place. we need to meet the feasibility condition first.

upvoted 3 times

⊟ **Colmenarez** 3 months, 2 weeks ago

AD of FS (409600) is not lower than FD of the successor (409600)

upvoted 1 times

⊟ **Brand** 3 months, 2 weeks ago

It looks to me that variance = 3 would do the trick as well.

upvoted 2 times

Question #492 *Topic 1*

Refer to the exhibit.

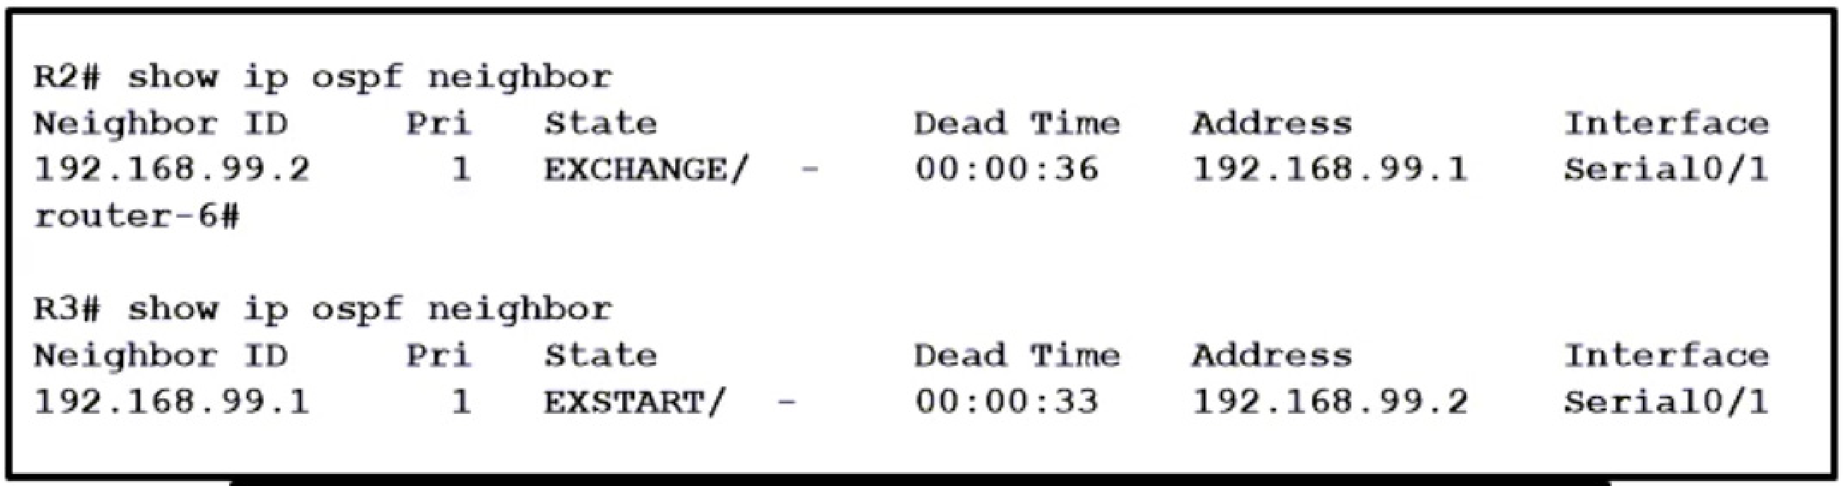

```
R2# show ip ospf neighbor
Neighbor ID      Pri    State           Dead Time    Address        Interface
192.168.99.2       1    EXCHANGE/  -    00:00:36     192.168.99.1   Serial0/1
router-6#

R3# show ip ospf neighbor
Neighbor ID      Pri    State           Dead Time    Address        Interface
192.168.99.1       1    EXSTART/  -     00:00:33     192.168.99.2   Serial0/1
```

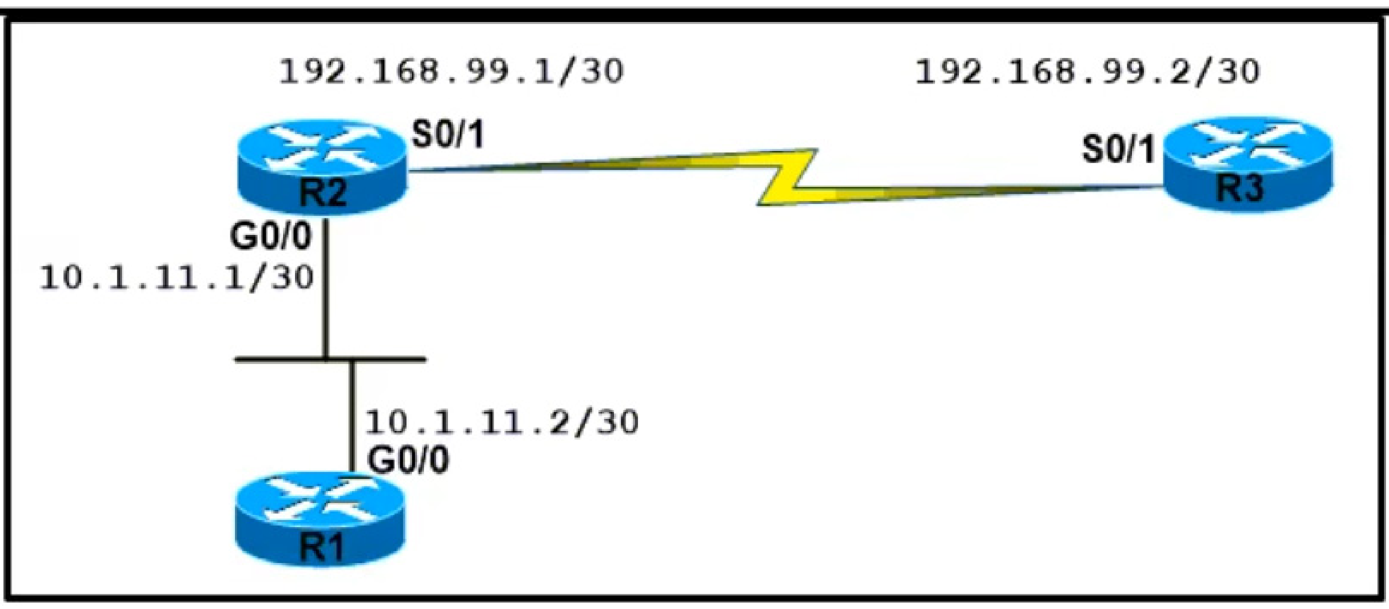

An OSPF neighbor relationship between R2 and R3 is showing stuck in EXCHANGE/EXSTART state. The neighbor is established between R1 and R2. The network engineer can ping from R2 to R3 and vice versa, but the neighbor is still down. Which action resolves the issue?

A. Administrative "shut then no shut" both router interfaces.

B. Enable OSPF on the interface, which is required.

C. Restore the Layer 2/Layer 3 connectivity issue in the ISP network.

D. Match MTU on both router interfaces or ignore MTU.

Correct Answer: *D*

seal2 3 months, 2 weeks ago

Selected Answer: D

A. Would do nothing to fix this except cause it to happen again
B. already enabled
C. clearly, they are already connected
D. this can fix the issue potentially

upvoted 1 times

Question #493 Topic 1

Refer to the exhibit.

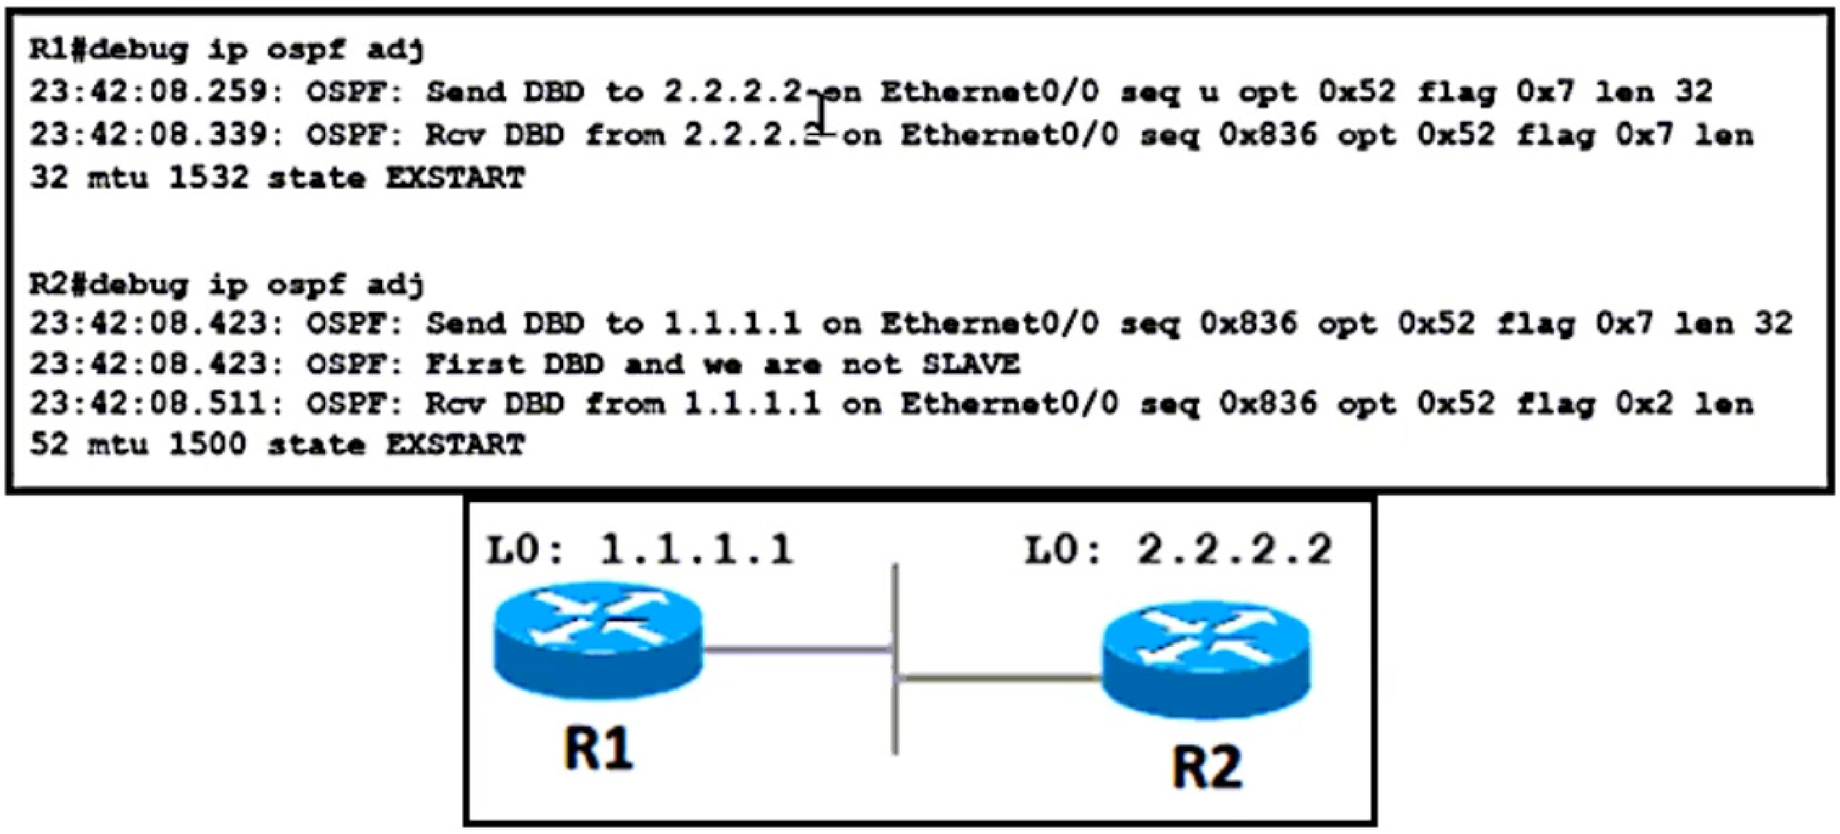

```
R1#debug ip ospf adj
23:42:08.259: OSPF: Send DBD to 2.2.2.2 on Ethernet0/0 seq u opt 0x52 flag 0x7 len 32
23:42:08.339: OSPF: Rcv DBD from 2.2.2.2 on Ethernet0/0 seq 0x836 opt 0x52 flag 0x7 len
32 mtu 1532 state EXSTART

R2#debug ip ospf adj
23:42:08.423: OSPF: Send DBD to 1.1.1.1 on Ethernet0/0 seq 0x836 opt 0x52 flag 0x7 len 32
23:42:08.423: OSPF: First DBD and we are not SLAVE
23:42:08.511: OSPF: Rcv DBD from 1.1.1.1 on Ethernet0/0 seq 0x836 opt 0x52 flag 0x2 len
52 mtu 1500 state EXSTART
```

R1 cannot establish a neighbor relationship with R2. Which action resolves the issue?

A. Configure the mtu ignore command on the interfaces of R1 and R2.

B. Configure the ip ospf network point-to-point command on the interfaces of R1 and R2.

C. Configure the neighbor 2.2.2.2 command on R1 under the OSPF process.

D. Configure the ip ospf network broadcast command on the interfaces of R1 and R2.

Correct Answer: *A*

seal2 3 months, 2 weeks ago

Selected Answer: A

A. would fix this, even if it is bad practice
B. would not fix MTU mismatch
C. looks like it's already configured?
D. would not fix anything

upvoted 1 times

Question #494 *Topic 1*

Refer to the exhibit.

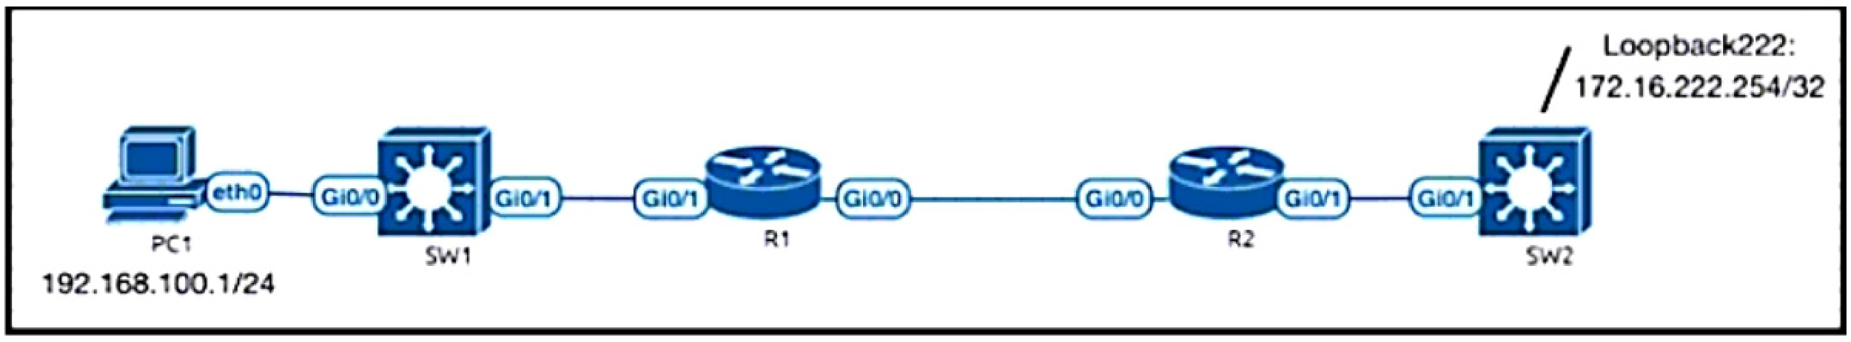

R2 can reach Loopback222, but R1, SW1, and PC1 cannot communicate with 172.16.222.254. R1 and R2 configurations are shown here:

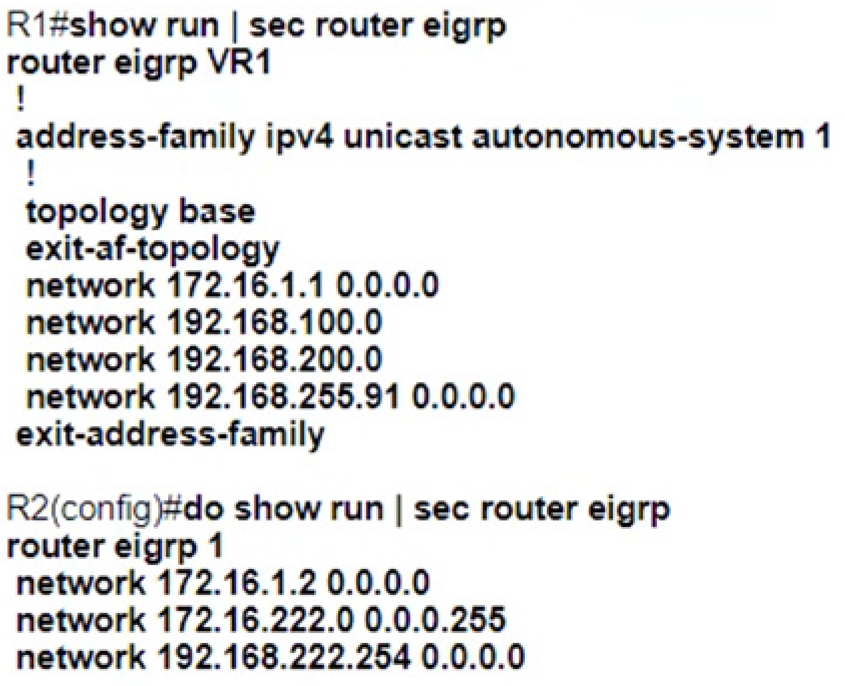

```
R1#show run | sec router eigrp
router eigrp VR1
 !
 address-family ipv4 unicast autonomous-system 1
  !
  topology base
  exit-af-topology
  network 172.16.1.1 0.0.0.0
  network 192.168.100.0
  network 192.168.200.0
  network 192.168.255.91 0.0.0.0
 exit-address-family

R2(config)#do show run | sec router eigrp
router eigrp 1
 network 172.16.1.2 0.0.0.0
 network 172.16.222.0 0.0.0.255
 network 192.168.222.254 0.0.0.0
```

Which EIGRP configuration command resolves the issue?

A. R1(config-router)# redistribute static

B. R2(config-router)# redistribute static

C. R1(config-router)# network 172.16.222.254 0.0.0.0

D. R1(config-router)# network 172.16.222.254 255.255.255.255

Correct Answer: *C*

seal2 Highly Voted 3 months, 2 weeks ago

Selected Answer: B

A. we don't need to redistribute on this router, it already doesn't have the route needed
B. we DO need to redistribute 172.16.222.0 0.0.0.255 in eigrp for it to work
C. without doing this on the eigrp instance it wouldnt be valid - this is my understanding from it being labeled (config-router)
D. same as C

upvoted 5 times

AlexInShort12 Most Recent 4 days, 7 hours ago

Selected Answer: B

Missing information, by expecting that their is a static route on R2, B is good.

upvoted 1 times

chris110 3 months, 1 week ago

Selected Answer: C

R1(config-router)# network 172.16.222.254 0.0.0.0

upvoted 1 times

chris110 3 months, 1 week ago

It is B

upvoted 1 times

SIMULATION

-

Guidelines

-

This is a lab item in which tasks will be performed on virtual devices.

• Refer to the Tasks tab to view the tasks for this lab item.
• Refer to the Topology tab to access the device console(s) and perform the tasks.
• Console access is available for all required devices by clicking the device icon or using the tab(s) above the console window.
• All necessary preconfigurations have been applied.
• Do not change the enable password or hostname for any device.
• Save your configurations to NVRAM before moving to the next item.
• Click Next at the bottom of the screen to submit this lab and move to the next question.
• When Next is clicked, the lab closes and cannot be reopened.

Topology

-

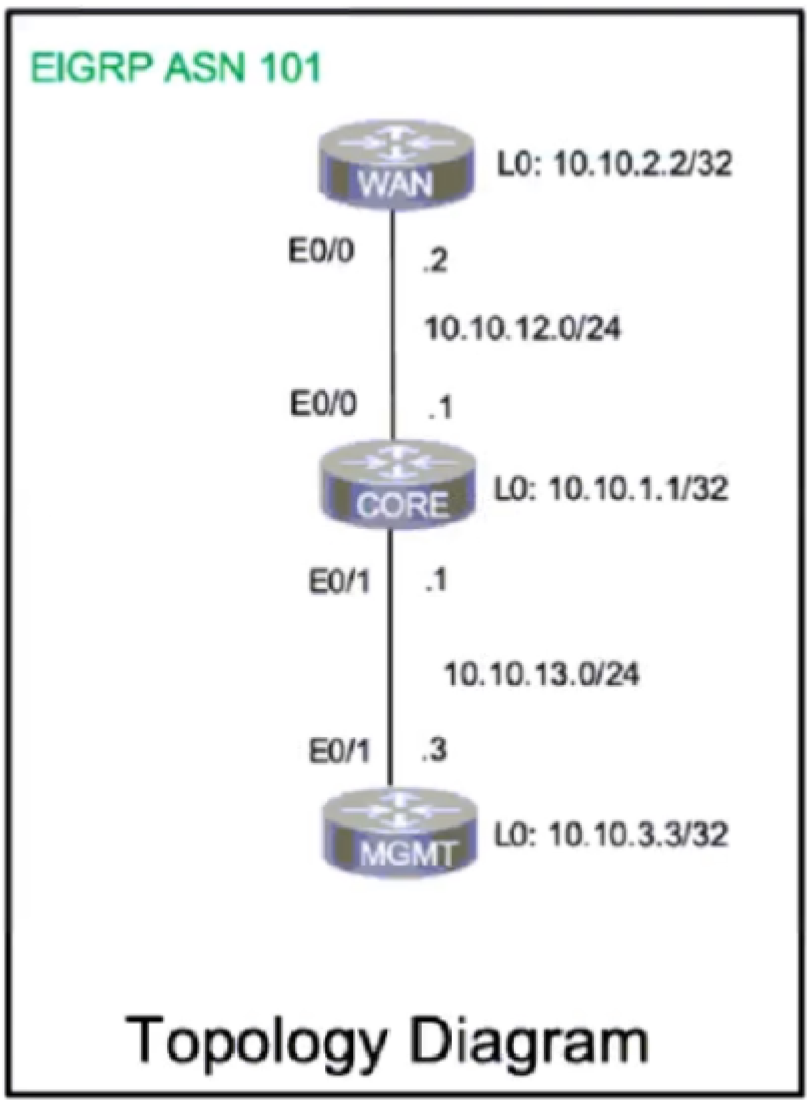

Tasks

-

A network is configured with CoPP to protect the CORE router route processor for stability and DDoS protection. As a company policy, a class named class-default is preconfigured and must not be modified or deleted. Troubleshoot CoPP to resolve the issues introduced during the maintenance window to ensure that:

1. Dynamic routing policies are under CoPP-CRITICAL and are allowed only from the 10.10.x.x range.

2. Telnet, SSH, and ping are under CoPP-IMPORTANT and are allowed strictly to/from 10.10.x.x to the CORE router (Hint: you can verify using

Loopback1).

3. All devices ping (UDP) any CORE router interface successfully to/from the 10.10.x.x range and do not allow any other IP address.

4. All devices run a successful traceroute (UDP) to any interface on the CORE router to/from the 10.10.x.x range, are under CoPP-NORMAL, and do not allow any other IP address traceroute is to be under CoPP-NORMAL (Hint: Traceroute port range 33434 33464).

Correct Answer:

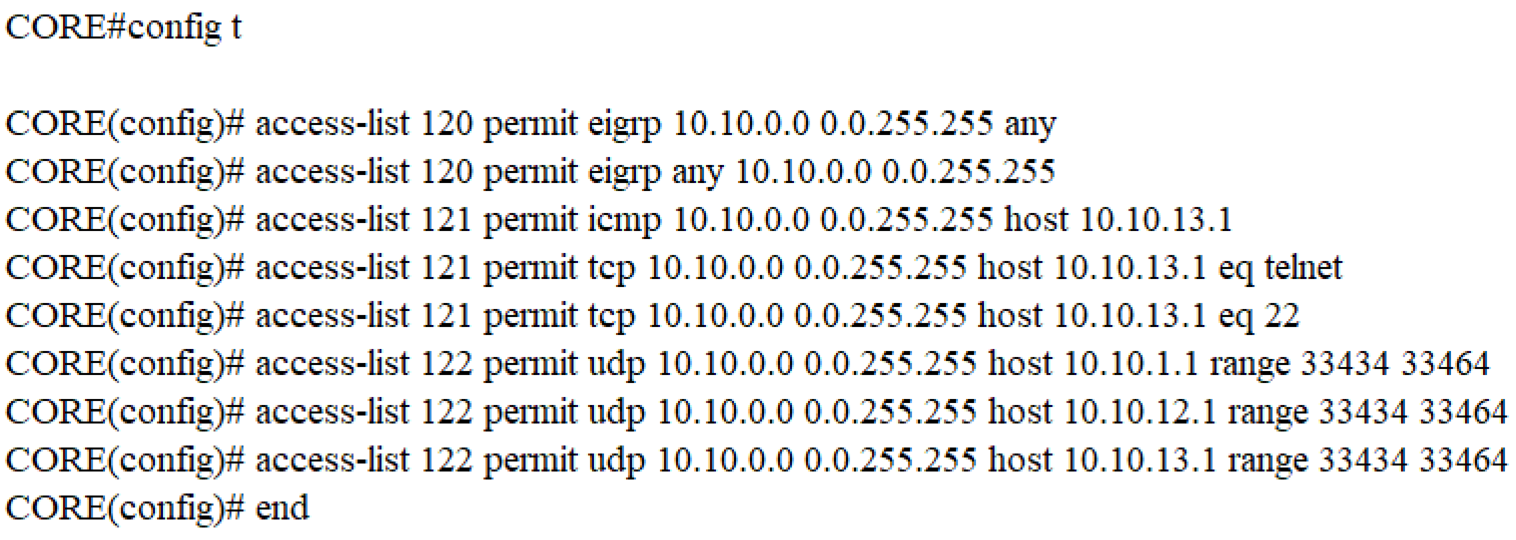

```
CORE#config t

CORE(config)# access-list 120 permit eigrp 10.10.0.0 0.0.255.255 any
CORE(config)# access-list 120 permit eigrp any 10.10.0.0 0.0.255.255
CORE(config)# access-list 121 permit icmp 10.10.0.0 0.0.255.255 host 10.10.13.1
CORE(config)# access-list 121 permit tcp 10.10.0.0 0.0.255.255 host 10.10.13.1 eq telnet
CORE(config)# access-list 121 permit tcp 10.10.0.0 0.0.255.255 host 10.10.13.1 eq 22
CORE(config)# access-list 122 permit udp 10.10.0.0 0.0.255.255 host 10.10.1.1 range 33434 33464
CORE(config)# access-list 122 permit udp 10.10.0.0 0.0.255.255 host 10.10.12.1 range 33434 33464
CORE(config)# access-list 122 permit udp 10.10.0.0 0.0.255.255 host 10.10.13.1 range 33434 33464
CORE(config)# end
```

⊟ **T_Cos** 1 day, 1 hour ago

The statement implies that you do not need to configure the policy, class-map or apply it to the control plane. Does anyone agree with me?

```
CORE(config)#ip access-list extended COPP-CRITICAL
(...-ext-nacl)#permit eigrp 10.10.0.0 0.0.255.255 any
(...-ext-nacl)#permit eigrp any 10.10.0.0 0.0.255.255
(...-ext-nacl)#permit eigrp any host 244.0.0.10

CORE(config)#ip access-list extended COPP-IMPORTANT
(...-ext-nacl)#permit icmp 10.10.0.0 0.0.255.255 host 10.10.13.1
(...-ext-nacl)#permit tcp 10.10.0.0 0.0.255.255 host 10.10.13.1 eq telnet
(...-ext-nacl)#permit tcp 10.10.0.0 0.0.255.255 host 10.10.13.1 eq 22

CORE(config)#ip access-list extended COPP-NORMAL
(...-ext-nacl)#permit udp 10.10.0.0 0.0.255.255 host 10.10.1.1 range 33434 33464
(...-ext-nacl)#permit udp 10.10.0.0 0.0.255.255 host 10.10.12.1 range 33434 33464
(...-ext-nacl)#permit udp 10.10.0.0 0.0.255.255 host 10.10.13.1 range 33434 33464
```

upvoted 1 times

⊟ **DeWalt95** 2 days, 11 hours ago

Anyone had this for real? The question and solution implies COPP policies/class-maps are setup and you just need to configure the ACLs? Also the wording suggests the ACLs are also applied to the interfaces?

upvoted 1 times

SIMULATION

-

Guidelines

-

This is a lab item in which tasks will be performed on virtual devices.

- Refer to the Tasks tab to view the tasks for this lab item.
- Refer to the Topology tab to access the device console(s) and perform the tasks.
- Console access is available for all required devices by clicking the device icon or using the tab(s) above the console window.
- All necessary preconfigurations have been applied.
- Do not change the enable password or hostname for any device.
- Save your configurations to NVRAM before moving to the next item.
- Click Next at the bottom of the screen to submit this lab and move to the next question.
- When Next is clicked, the lab closes and cannot be reopened.

Topology

-

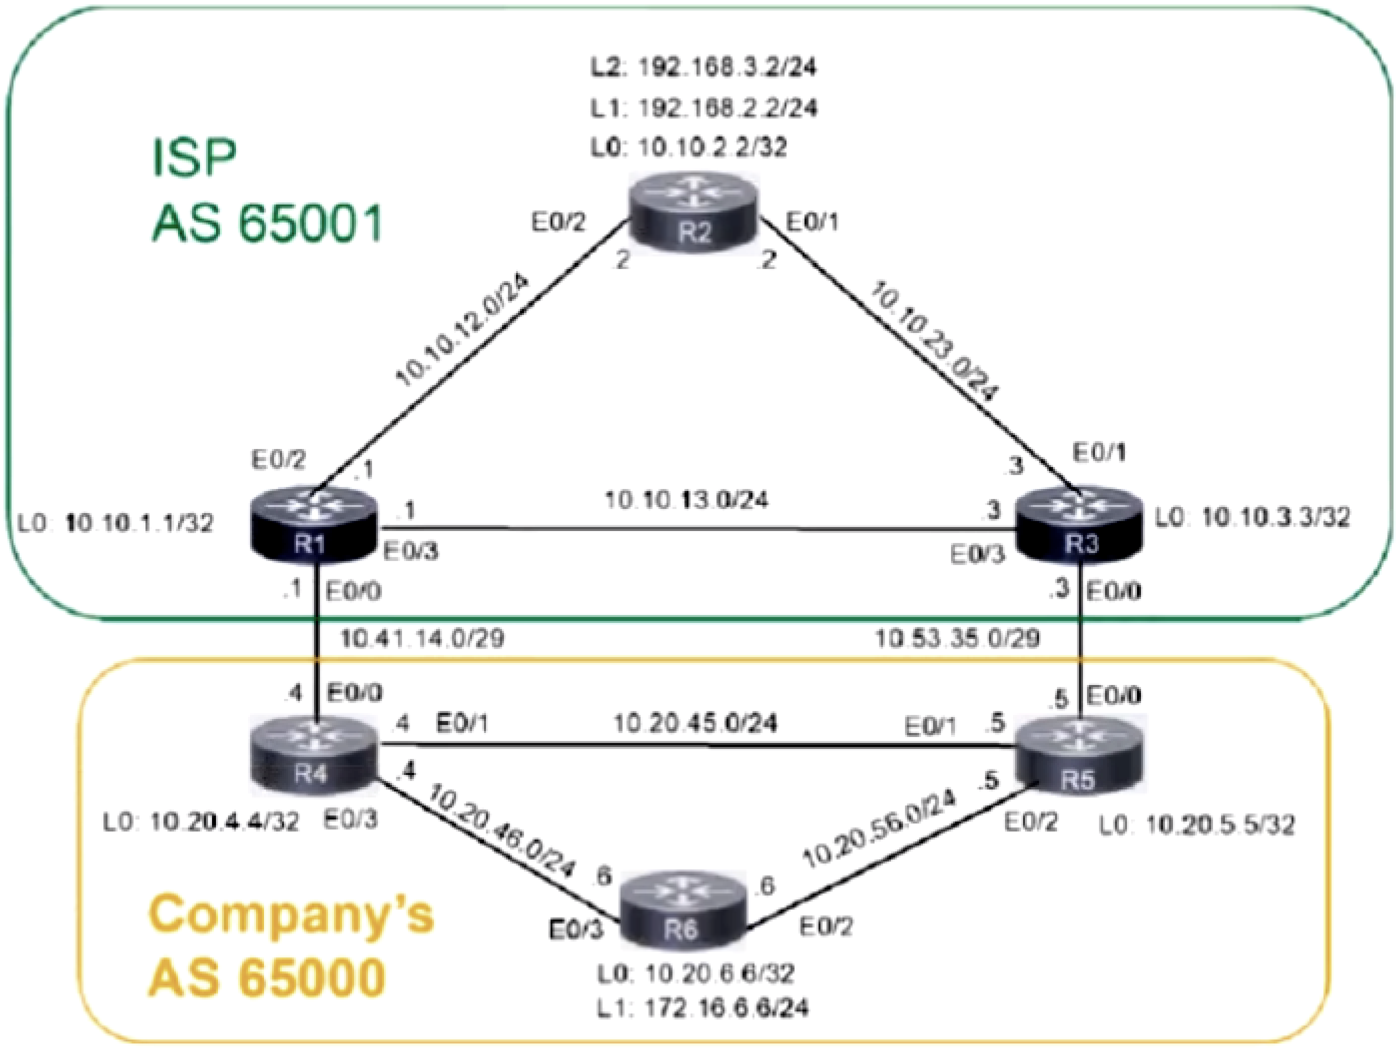

Tasks

-

A company is connected to an ISP and some of the networks between the ISP and the company are not reachable. Troubleshoot and resolve the issues to achieve these goals:

1. A single /16 is advertised for all infrastructure-connected interfaces that belong to the 10.20.x.x network using BGP network commands from border routers connected to the ISP. Configuration modification is allowed in R4 and R5 to achieve the results. Do not use the BGP aggregate command.

2. R6 receives the ISP R2 Loopback2 from R4 and receives a summary address for both Loopbacks of ISP R2 from R4 or R5. Use BGP attribute

local-preference, add , for example, for R6, use “default+6 = value to be used”. Use the existing prefix lists or route maps with the sequence numbering starting at 10 and added in increments of 10.

3. R6 receives the ISP R2 Loopback1 from R5 and receives a summary address for both Loopbacks of ISP R2 from R4 or R5 using the same guidelines.

4. R6 advertises its Loopback1 /24 address through BGP.

Correct Answer:

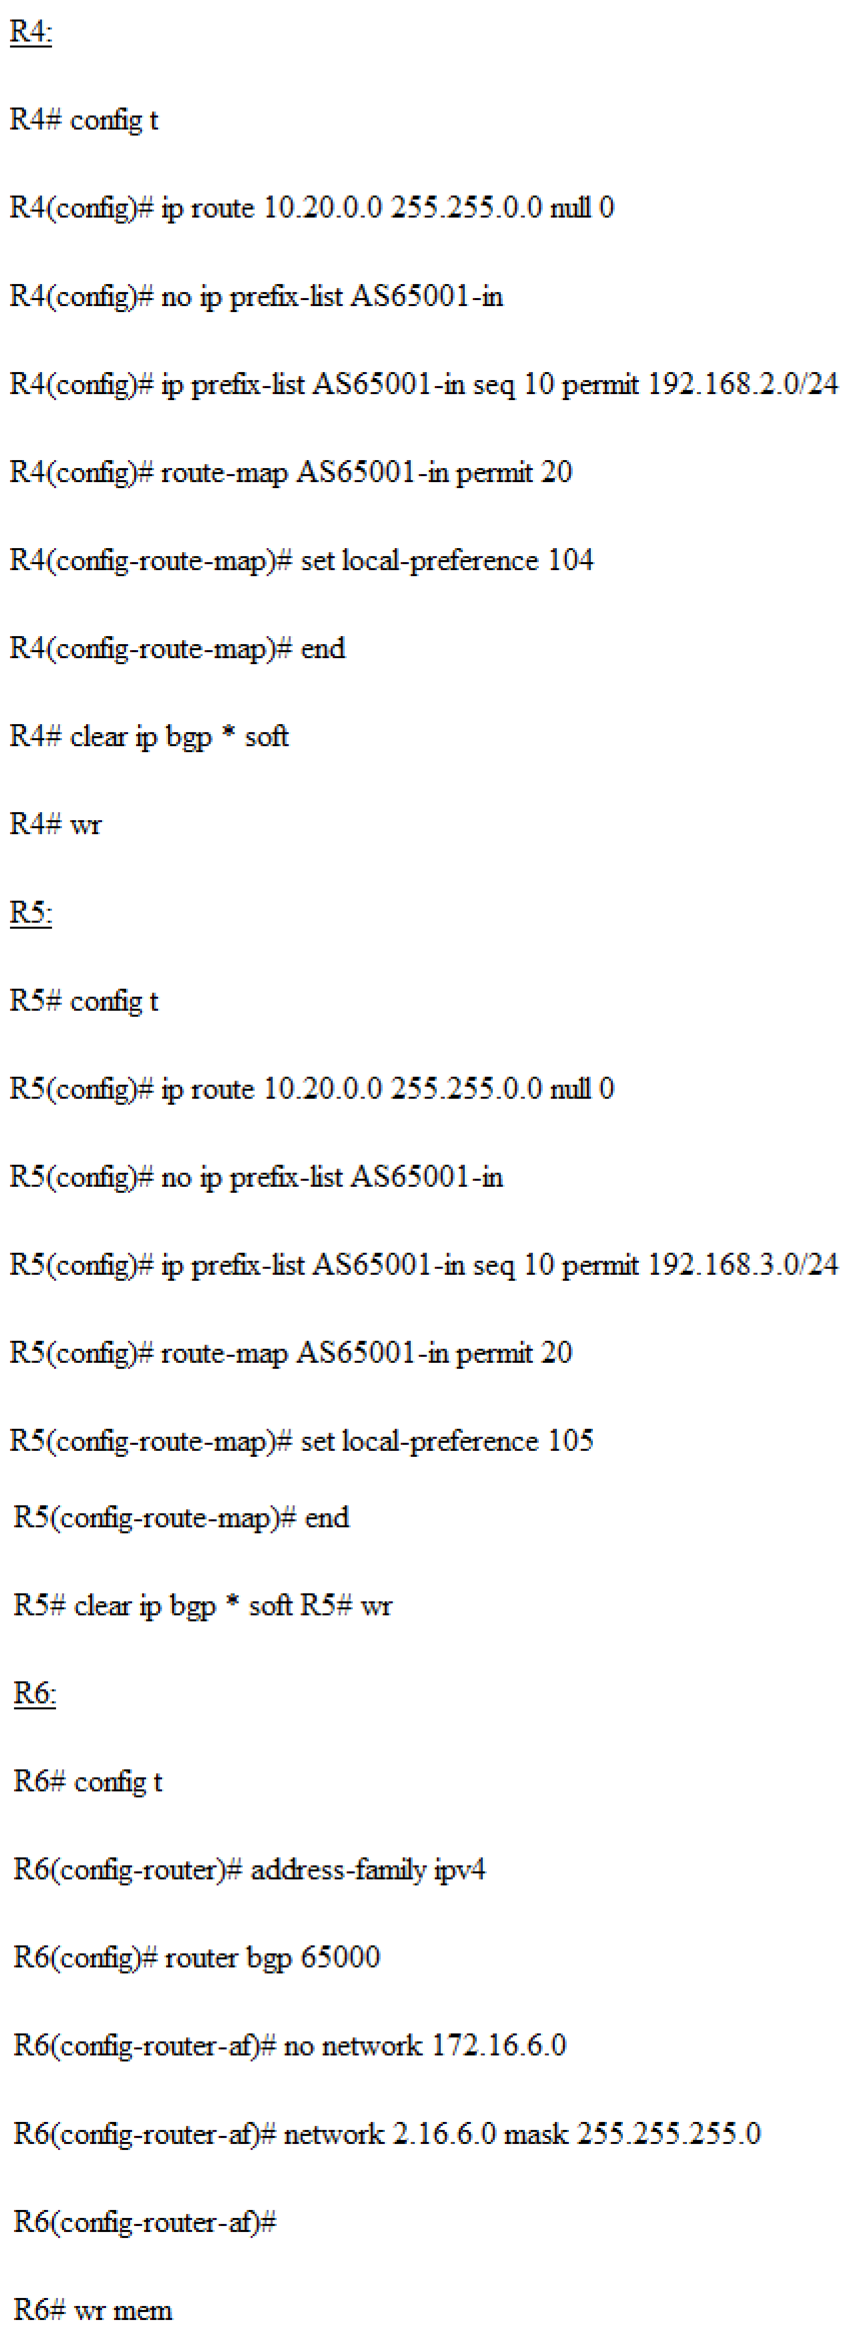

R4:

```
R4# config t

R4(config)# ip route 10.20.0.0 255.255.0.0 null 0

R4(config)# no ip prefix-list AS65001-in

R4(config)# ip prefix-list AS65001-in seq 10 permit 192.168.2.0/24

R4(config)# route-map AS65001-in permit 20

R4(config-route-map)# set local-preference 104

R4(config-route-map)# end

R4# clear ip bgp * soft

R4# wr
```

R5:

```
R5# config t

R5(config)# ip route 10.20.0.0 255.255.0.0 null 0

R5(config)# no ip prefix-list AS65001-in

R5(config)# ip prefix-list AS65001-in seq 10 permit 192.168.3.0/24

R5(config)# route-map AS65001-in permit 20

R5(config-route-map)# set local-preference 105

R5(config-route-map)# end

R5# clear ip bgp * soft R5# wr
```

R6:

```
R6# config t

R6(config-router)# address-family ipv4

R6(config)# router bgp 65000

R6(config-router-af)# no network 172.16.6.0

R6(config-router-af)# network 2.16.6.0 mask 255.255.255.0

R6(config-router-af)#

R6# wr mem
```

⊟ **DeWalt95** 3 days, 6 hours ago

Labbed as best I could without seeing the actual configs.

Part 1
Router 4+5
Add static route to 10.20.0.0 255.255.0.0 null0
Advertise to BGP with network statement
Part 2
Router 4 and 5.

Match relevant R2 loopback using a standard access list.
Create route map (question implies modification of existing). Permit 10 match access list and set LP to 100+routerID. Permit 20 match all/permit all (blank).

Part3
Advertise Loopback on R6 using network statement.
upvoted 1 times

Ghauri777 1 month ago

```
R4
router bgp 65000
bgp log-neighbor-changes
network 10.20.0.0 mask 255.255.0.0
aggregate-address 192.168.0.0 255.255.252.0 as-set
neighbor 10.10.1.1 remote-as 65001
neighbor 10.10.1.1 ebgp-multihop 2
neighbor 10.10.1.1 update-source Loopback0
neighbor 10.10.1.1 route-map R2-L2 in


ip prefix-list 1 seq 5 permit 192.168.3.0/24
match ip address prefix-list 1


route-map R2-L2 permit 10
match ip address prefix-list 1
set local-preference 104
route-map R2-L2 permit 20
-----------------------------------

R5
router bgp 65000
bgp log-neighbor-changes
network 10.20.0.0 mask 255.255.0.0
aggregate-address 192.168.0.0 255.255.252.0 as-set
neighbor 10.10.3.3 remote-as 65001
neighbor 10.10.3.3 ebgp-multihop 2
neighbor 10.10.3.3 update-source Loopback0
neighbor 10.10.3.3 route-map R2-L1 in

ip prefix-list 1 seq 5 permit 192.168.2.0/24
match ip address prefix-list 1
set local-preference 106

route-map R2-L1 permit 10
match ip address prefix-list 1
set local-preference 106
route-map R2-L1 permit 20
---------------

R6
router bgp 65000
bgp log-neighbor-changes
network 172.16.6.0 mask 255.255.255.0
```

upvoted 1 times

DeWalt95 2 weeks, 4 days ago

Its says you cannot use the aggregate command.

As the published solution states - add a null static route and then advertise into BGP
upvoted 1 times

aqwsdfghjklp 3 weeks, 3 days ago

Is this a config?
upvoted 1 times

aqwsdfghjklp 3 weeks, 5 days ago

I don't know why you need the "as-set" command.
upvoted 1 times

Question #497

Topic 1

How does BFD protocol work?

A. When BFD declares a failure on the primary IGP path, the router on the peer router chooses to use the secondary path.

B. BFD operates on the route processor module and impacts the route processor CPU utilization.

C. BFD control packets are sent via UDP port 3784 to the destination router.

D. BFD echo packets are sent to the same source IP and different destination IP with TCP port of 3786.

Correct Answer: *C*

Rman0059 1 month, 2 weeks ago

Selected Answer: C

C - BFD uses port UDP 3784

upvoted 1 times

Question #498 *Topic 1*

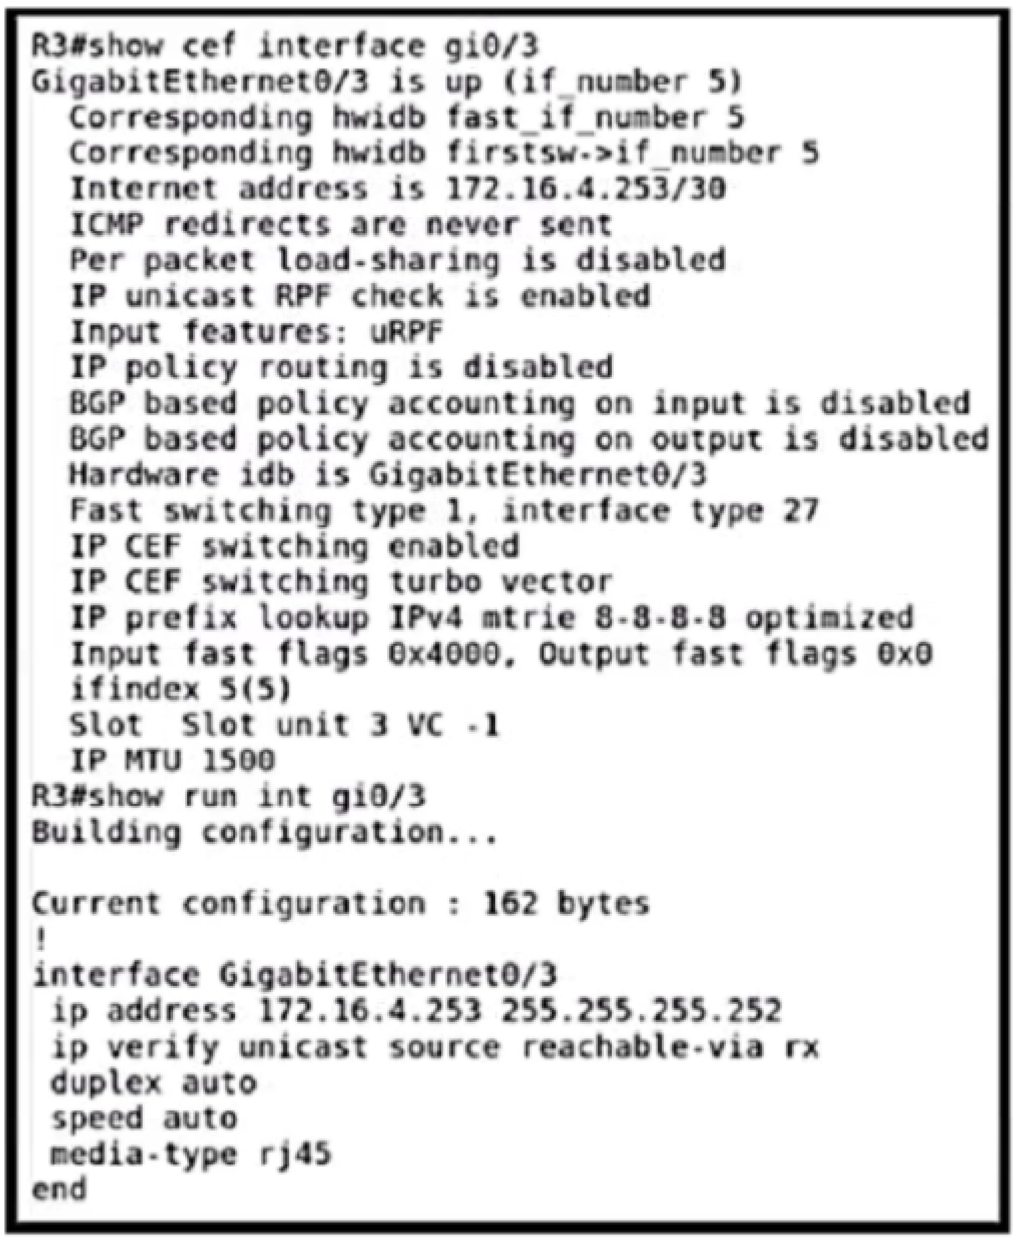

```
R3#show cef interface gi0/3
GigabitEthernet0/3 is up (if_number 5)
  Corresponding hwidb fast_if_number 5
  Corresponding hwidb firstsw->if_number 5
  Internet address is 172.16.4.253/30
  ICMP redirects are never sent
  Per packet load-sharing is disabled
  IP unicast RPF check is enabled
  Input features: uRPF
  IP policy routing is disabled
  BGP based policy accounting on input is disabled
  BGP based policy accounting on output is disabled
  Hardware idb is GigabitEthernet0/3
  Fast switching type 1, interface type 27
  IP CEF switching enabled
  IP CEF switching turbo vector
  IP prefix lookup IPv4 mtrie 8-8-8-8 optimized
  Input fast flags 0x4000, Output fast flags 0x0
  ifindex 5(5)
  Slot  Slot unit 3 VC -1
  IP MTU 1500
R3#show run int gi0/3
Building configuration...

Current configuration : 162 bytes
!
interface GigabitEthernet0/3
 ip address 172.16.4.253 255.255.255.252
 ip verify unicast source reachable-via rx
 duplex auto
 speed auto
 media-type rj45
end
```

Refer to the exhibit. An engineer implements uRPF to increase security and stop incoming spoofed IP packets. Same asymmetrically routed packets are also blocked after the configuration. Which command resolves the issue?

A. ip verify unicast source reachable-via any

B. ip verify unicast source reachable-via rx

C. ip verify unicast reverse-path

D. ip verify unicast reverse-path any

Correct Answer: *A*

Rman0059 1 month, 2 weeks ago

Selected Answer: A

A is correct. Need to issue the loose RPF commands

upvoted 1 times

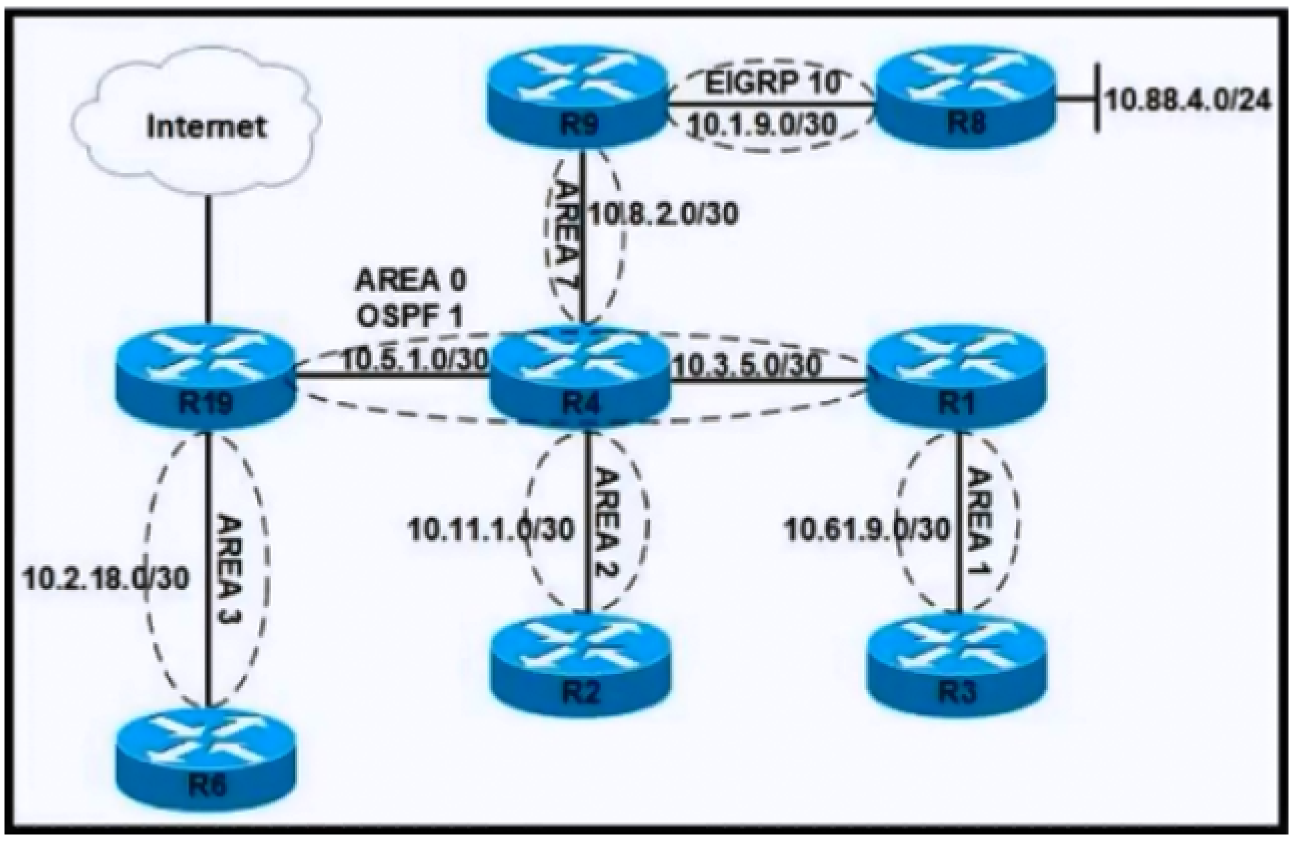

Refer to the exhibit. After an engineer modified the configuration for area 7 to permit type 1, 2, and 7 LSAs only, users connected to router R9 reported that they could no longer access the internet. Which configuration restores internet access to users on R9 and permits only LSA type 1, 2, and 7?

A. R4#
```
router ospf 1
area 0 nssa default-information-originate
network 10.5.1.0 0.0.0.3 area 0
network 10.8.2.0 0.0.0.3 area 7

R9#
router ospf 1
area 7 nssa
redistribute eigrp 10 subnets
network 10.8.2.0 0.0.0.3 area 7
```

B. R4#
```
router ospf 1
area 7 nssa no-summary
network 10.5.1.0 0.0.0.3 area 0
network 10.8.2.0 0.0.0.3 area 7

R9#
router ospf 1
area 7 nssa
redistribute eigrp 10 subnets
network 10.8.2.0 0.0.0.3 area 7
```

C. R4#
```
router ospf 1
area 7 nssa
network 10.5.1.0 0.0.0.3 area 0
network 10.8.2.0 0.0.0.3 area 7

R9#
router ospf 1
area 7 nssa
redistribute eigrp 10 subnets
network 10.8.2.0 0.0.0.3 area 7
```

D. R4#

```
router ospf 1
area 0 area 7 stub no-summary
network 10.5.1.0 0.0.0.3 area 0
network 10.8.2.0 0.0.0.3 area 7
```

R9#

```
router ospf 1
area 7 stub
redistribute eigrp 10 subnets
network 10.8.2.0 0.0.0.3 area 7
```

Correct Answer: *A*

AlexInShort12 4 days, 6 hours ago

Selected Answer: B

Missing configuration in B answer
area 7 nssa -> no-summary

upvoted 1 times

b8os5h 1 month, 1 week ago

Selected Answer: B

A.NSSA
(If the command is "area 7 nssa default-information-originate")
B.Totally NSSA
NSSA blocks LSA types 4 and 5
Totally NSSA blocks LSA types 3, 4, and 5
In NSSA, ABR does not automatically advertise a default route.
Therefore, the "default-informationoriginate" keyword must be configured while configuring NSSA.
In the case of Totally NSSA, there is no need to configure "default-information originate" to advertise the default route.

upvoted 1 times

Question #500 *Topic 1*

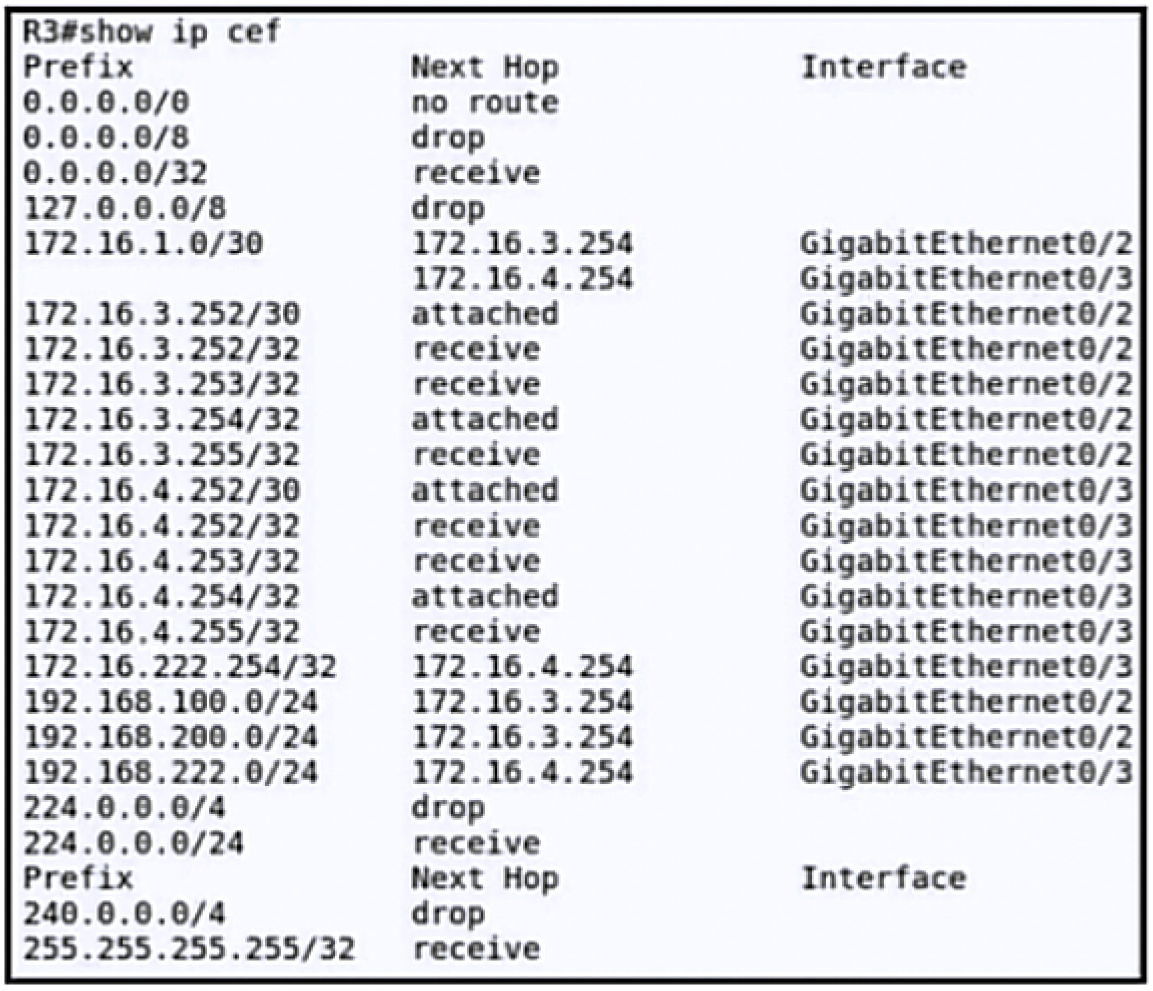

```
R3#show ip cef
Prefix              Next Hop            Interface
0.0.0.0/0           no route
0.0.0.0/8           drop
0.0.0.0/32          receive
127.0.0.0/8         drop
172.16.1.0/30       172.16.3.254        GigabitEthernet0/2
                    172.16.4.254        GigabitEthernet0/3
172.16.3.252/30     attached            GigabitEthernet0/2
172.16.3.252/32     receive             GigabitEthernet0/2
172.16.3.253/32     receive             GigabitEthernet0/2
172.16.3.254/32     attached            GigabitEthernet0/2
172.16.3.255/32     receive             GigabitEthernet0/2
172.16.4.252/30     attached            GigabitEthernet0/3
172.16.4.252/32     receive             GigabitEthernet0/3
172.16.4.253/32     receive             GigabitEthernet0/3
172.16.4.254/32     attached            GigabitEthernet0/3
172.16.4.255/32     receive             GigabitEthernet0/3
172.16.222.254/32   172.16.4.254        GigabitEthernet0/3
192.168.100.0/24    172.16.3.254        GigabitEthernet0/2
192.168.200.0/24    172.16.3.254        GigabitEthernet0/2
192.168.222.0/24    172.16.4.254        GigabitEthernet0/3
224.0.0.0/4         drop
224.0.0.0/24        receive
Prefix              Next Hop            Interface
240.0.0.0/4         drop
255.255.255.255/32  receive
```

Refer to the exhibit. An engineer recently implemented uRPF by configuring the ip verify unicast source reachable-via rx command on interface gi0/3. The engineer noticed right after implementing uRPF that an inbound packet on the gi0/3 interface with a source address of 172.16.3.251 was dropped. Which action resolves the issue?

A. Configure uRPF loose mode to forward the packet.

B. Permit the 172.16.3.251 in the inbound ACL on interface gi0/3.

C. Remove inbound ACL from the interface gi0/3 to allow 172.16.3.251.

D. Configure uRPF strict mode to forward the packet.

Correct Answer: *A*

Rman0059 1 month, 2 weeks ago

Selected Answer: A

A is correct. Need to enable loose mode

upvoted 1 times

Question #501 *Topic 1*

Which IPv6 feature enables a device to reject traffic when it is originated from an address that is not stored in the device binding table?

A. IPv6 Source Guard

B. IPv6 DAD Proxy

C. IPv6 RA Guard

D. IPv6 Snooping

Correct Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #502 *Topic 1*

Which two protocols are used by a P router to transfer VPN traffic between PE routers in an MPLS network? (Choose two.)

A. LDP

B. RSVP

C. MP-BGP

D. BGP

E. OSPF

Correct Answer: *AE*

night_wolf_in Highly Voted 1 month, 1 week ago

Selected Answer: AB

P to PE uses LDP and RSVP.

upvoted 6 times

Tedmus Most Recent 1 week, 6 days ago

Selected Answer: AB

To enable MPLS on all routers in the core, you must configure a label distribution protocol. You can use either of the following as a label distribution protocol:

• MPLS Label Distribution Protocol (LDP).

• MPLS Traffic Engineering Resource Reservation Protocol (RSVP)

upvoted 1 times

Question #503 *Topic 1*

Which feature is used by LDP in the forwarding path within the MPLS cloud?

A. TDP

B. TTL

C. LSP

D. IP forwarding

Correct Answer: *B*

Tedmus 1 week, 6 days ago

Selected Answer: D

Reference:
https://www.cisco.com/en/US/docs/ios-xml/ios/mp_ldp/configuration/15-2s/mp-ldp-overview.html

...method of label distribution is also called hop-by-hop forwarding. With IP forwarding...

We are searching for a "Feature".
LSP is the establish a label-switched path.

upvoted 1 times

T_Cos 2 weeks, 1 day ago

Option "C"

upvoted 1 times

Normanby 3 weeks, 3 days ago

Selected Answer: C

Answer is C (LSP)

upvoted 4 times

Wanwan 4 weeks, 1 day ago

Answer is C (LSP)

upvoted 2 times

Question #504 *Topic 1*

Which two features are required for MPLS forwarding on which types of routers? (Choose two.)

A. MPLS on PE and core routers

B. LDP on PE and core routers

C. MPLS on CE and core routers

D. LDP on PE and CE routers

E. CEF on PE and CE routers

Correct Answer: *AB*

Question #505 *Topic 1*

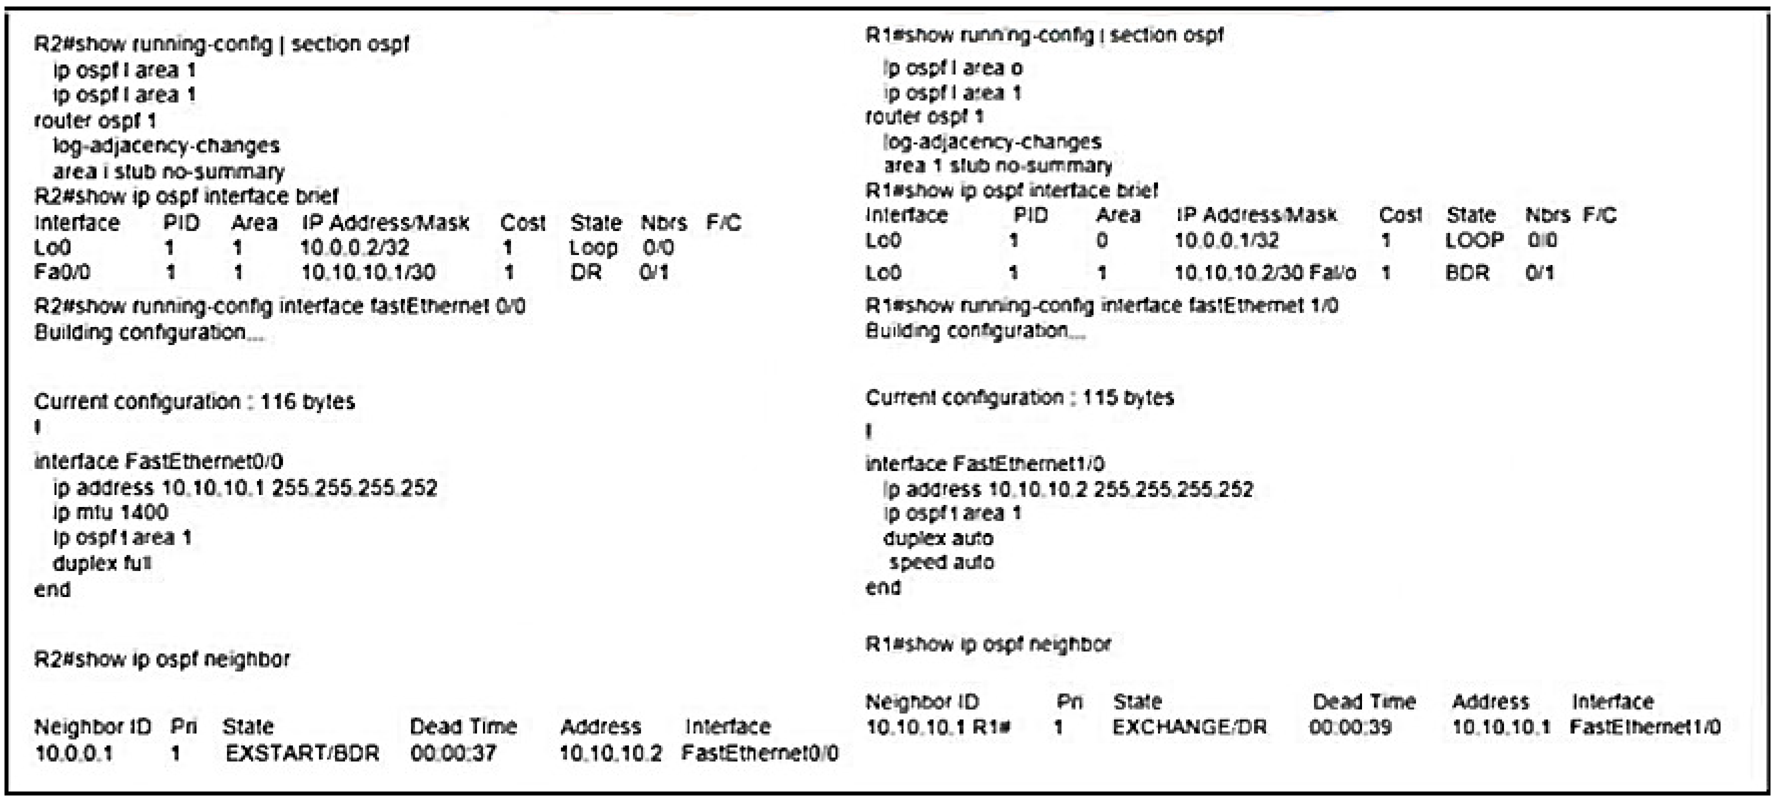

```
R2#show running-config | section ospf
  ip ospf I area 1
  ip ospf I area 1
router ospf 1
  log-adjacency-changes
  area i stub no-summary
R2#show ip ospf interface brief
Interface    PID   Area   IP Address/Mask    Cost   State   Nbrs   F/C
Lo0          1     1      10.0.0.2/32        1      Loop    0/0
Fa0/0        1     1      10.10.10.1/30      1      DR      0/1
R2#show running-config interface fastEthernet 0/0
Building configuration...


Current configuration : 116 bytes
!
interface FastEthernet0/0
  ip address 10.10.10.1 255.255.255.252
  ip mtu 1400
  ip ospf 1 area 1
  duplex full
end


R2#show ip ospf neighbor


Neighbor ID   Pri   State          Dead Time    Address      Interface
10.0.0.1      1     EXSTART/BDR    00:00:37     10.10.10.2   FastEthernet0/0
```

```
R1#show running-config | section ospf
  Ip ospf I area o
  ip ospf I area 1
router ospf 1
  log-adjacency-changes
  area 1 stub no-summary
R1#show ip ospf interface brief
Interface      PID     Area    IP Address/Mask          Cost   State   Nbrs  F/C
Lo0            1       0       10.0.0.1/32              1      LOOP    0/0
Lo0            1       1       10.10.10.2/30 Fa1/o      1      BDR     0/1
R1#show running-config interface fastEthernet 1/0
Building configuration...


Current configuration : 115 bytes
!
interface FastEthernet1/0
  Ip address 10.10.10.2 255.255.255.252
  ip ospf 1 area 1
  duplex auto
   speed auto
end

R1#show ip ospf neighbor

Neighbor ID        Pri   State            Dead Time    Address      Interface
10.10.10.1 R1#     1     EXCHANGE/DR      00:00:39     10.10.10.1   FastEthernet1/0
```

Refer to the exhibit. Which action restores OSPF adjacency between R1 and R2?

A. Change the IP MTU of R2 Fa0/0 to 1300.

B. Change the IP MTU of R1 Fa1/0 to 1500.

C. Change the IP MTU of R2 Fa0/0 to 1500.

D. Change the IP MTU of R1 Fa1/0 to 1300.

Correct Answer: *C*

Question #506 *Topic 1*

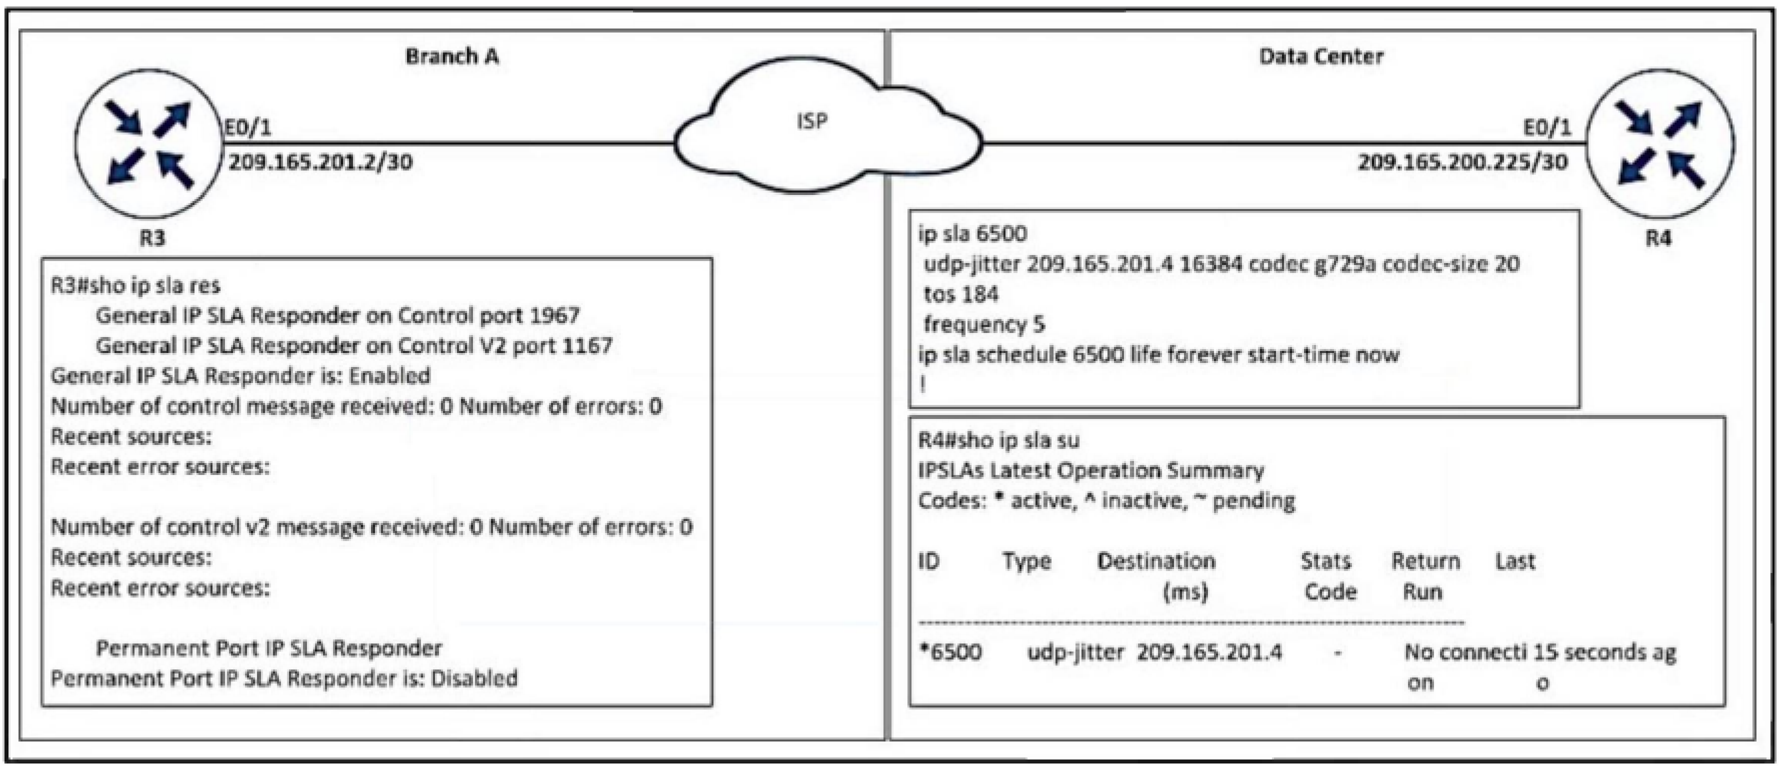

```
R3#sho ip sla res
      General IP SLA Responder on Control port 1967
      General IP SLA Responder on Control V2 port 1167
General IP SLA Responder is: Enabled
Number of control message received: 0 Number of errors: 0
Recent sources:
Recent error sources:

Number of control v2 message received: 0 Number of errors: 0
Recent sources:
Recent error sources:

      Permanent Port IP SLA Responder
Permanent Port IP SLA Responder is: Disabled
```

```
ip sla 6500
 udp-jitter 209.165.201.4 16384 codec g729a codec-size 20
 tos 184
 frequency 5
ip sla schedule 6500 life forever start-time now
!
```

```
R4#sho ip sla su
IPSLAs Latest Operation Summary
Codes: * active, ^ inactive, ~ pending

ID       Type       Destination        Stats    Return    Last
                          (ms)         Code      Run
------------------------------------------------------------------
*6500    udp-jitter  209.165.201.4       -      No connecti 15 seconds ag
                                                 on           o
```

Refer to the exhibit. Which action resolves the IP SLA for the UDP jitter problem between R4 and R3 Ethernet 0/1 IP addresses?

A. Delete and configure the ip sla 6500 command with R3 e0/1 IP address.

B. Configure the ip sla 6500 command with R3 e0/1 IP address.

C. Configure the ip sla responder command with R4 E0/1 IP address.

D. Delete and configure the ip sla responder command with R4 E0/1 IP address.

Correct Answer: *A*

Question #507 *Topic 1*

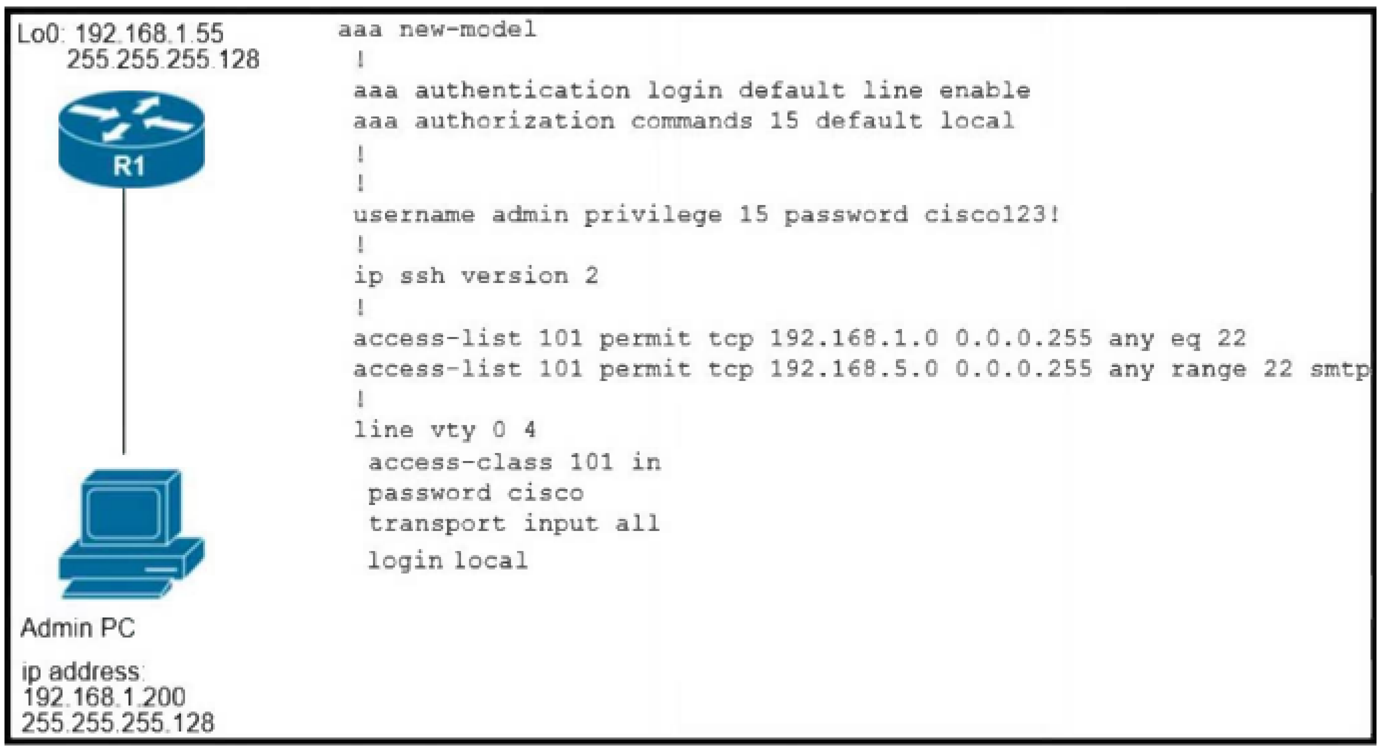

```
aaa new-model
 !
 aaa authentication login default line enable
 aaa authorization commands 15 default local
 !
 !
 username admin privilege 15 password cisco123!
 !
 ip ssh version 2
 !
 access-list 101 permit tcp 192.168.1.0 0.0.0.255 any eq 22
 access-list 101 permit tcp 192.168.5.0 0.0.0.255 any range 22 smtp
 !
 line vty 0 4
  access-class 101 in
  password cisco
  transport input all
  login local
```

Refer to the exhibit. An engineer configured user login based on authentication database on the router, but no one can log into the router. Which configuration resolves the issue?

A. aaa authentication login default local

B. aaa authorization network default local

C. aaa authentication login default enable

D. aaa authorization exec default local

Correct Answer: *A*

Question #508 *Topic 1*

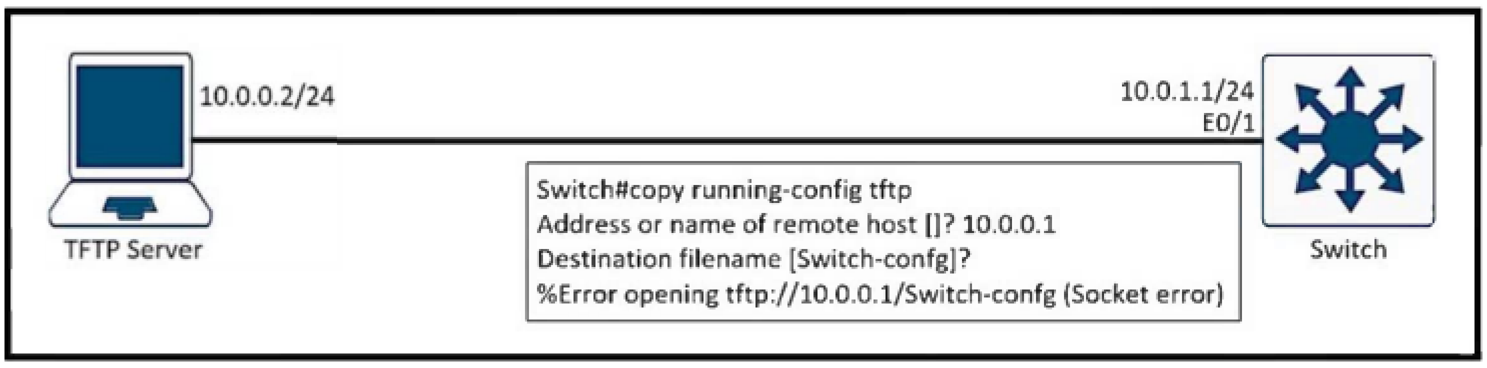

```
Switch#copy running-config tftp
Address or name of remote host []? 10.0.0.1
Destination filename [Switch-confg]?
%Error opening tftp://10.0.0.1/Switch-confg (Socket error)
```

Refer to the exhibit. Which action allows the engineer to successfully copy running-config to the TFTP server?

A. Add a route in the switch to the TFTP server.

B. Add the TFTP server configuration in the switch.

C. Use TFTP server IP address 10.0.1.1.

D. Use file name Switch-confg.txt.

Correct Answer: *C*

⊟ **AlexInShort12** 4 days, 6 hours ago

Missing information in the diagram.
They are not even on the same subnet.

upvoted 2 times

⊟ **DavideDL** 1 week, 1 day ago

In my opinion C makes sense if the ip addresses would be swapped...

upvoted 2 times

Question #509 Topic 1

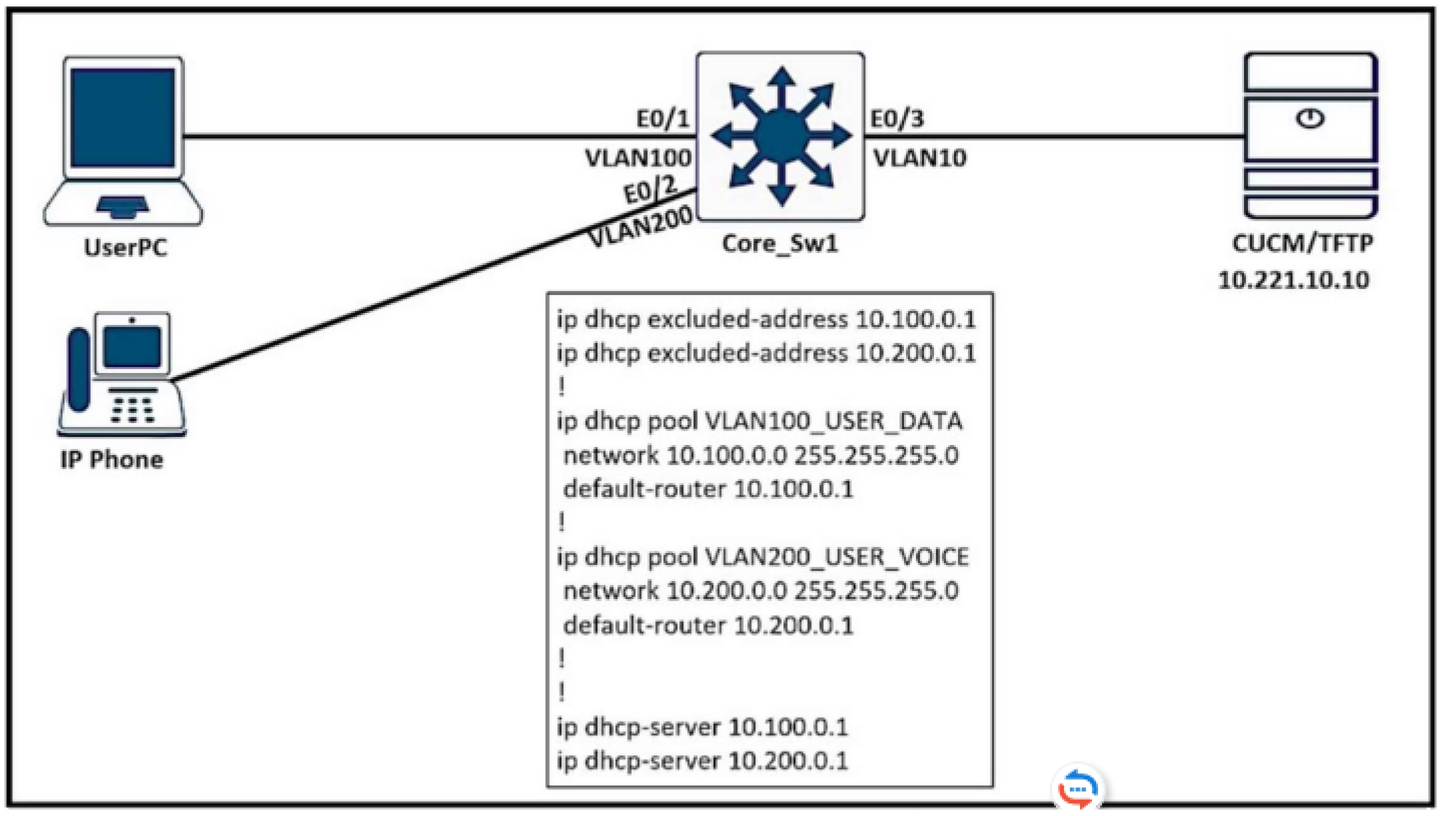

```
ip dhcp excluded-address 10.100.0.1
ip dhcp excluded-address 10.200.0.1
!
ip dhcp pool VLAN100_USER_DATA
 network 10.100.0.0 255.255.255.0
 default-router 10.100.0.1
!
ip dhcp pool VLAN200_USER_VOICE
 network 10.200.0.0 255.255.255.0
 default-router 10.200.0.1
!
!
ip dhcp-server 10.100.0.1
ip dhcp-server 10.200.0.1
```

Refer to the exhibit. UserPC receives the IP address but does not register to the call manager. Which command in ip dhcp pool VLAN200_USER_VOICE resolves the issue?

A. option 150 ip 10.221.10.10

B. option 15 ip 10.221.10.10

C. option 160 ip 10.221.10.10

D. option 117 ip 10.221.10.10

Correct Answer: *A*

DeWalt95 2 weeks, 3 days ago

Selected Answer: A

DHCP option 150 = TFTP server address

upvoted 2 times

Question #510 Topic 1

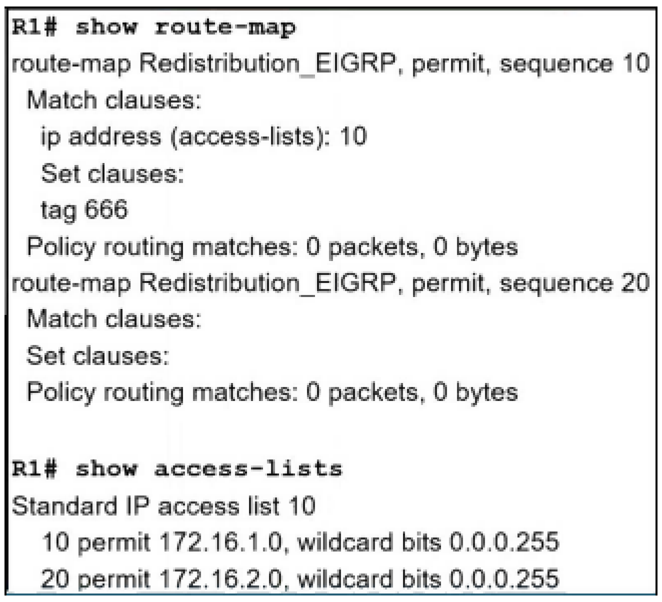

```
R1# show route-map
route-map Redistribution_EIGRP, permit, sequence 10
  Match clauses:
    ip address (access-lists): 10
    Set clauses:
    tag 666
  Policy routing matches: 0 packets, 0 bytes
route-map Redistribution_EIGRP, permit, sequence 20
  Match clauses:
  Set clauses:
  Policy routing matches: 0 packets, 0 bytes

R1# show access-lists
Standard IP access list 10
    10 permit 172.16.1.0, wildcard bits 0.0.0.255
    20 permit 172.16.2.0, wildcard bits 0.0.0.255
```

Refer to the exhibit. The router is redistributing a prefix 172.16.10.0/24 that should have been filtered. Which action resolves the issue?

A. Add the route in access-list 10.

B. Match the tag 666 for the route in the route map.

C. Remove route-map sequence 20.

D. Permit the route in route-map sequence 20.

Correct Answer: *C*

⊟ **asans** 1 week, 2 days ago

C is correct, Remove route-map sequence 20 and that denies redistributing any traffic that doesnt match acl 10

upvoted 1 times

⊟ **changer30** 2 weeks, 3 days ago

none of the answer have to do with the prefix 172.16.10.0/24, none of it involve deny/blocking. I'm really frustrated with the exam. good work cisco

upvoted 2 times

⊟ **DeWalt95** 2 weeks, 3 days ago

Selected Answer: C

All traffic permitted by sequence 20

upvoted 2 times

Question #511 *Topic 1*

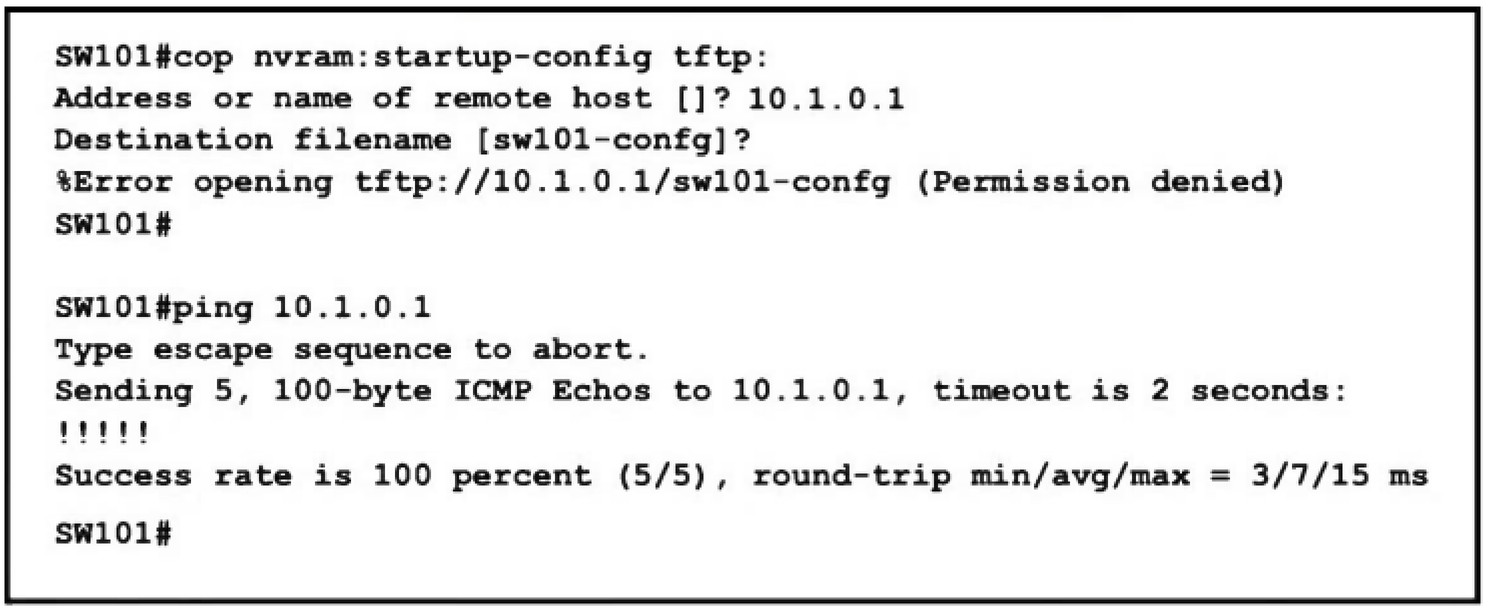

```
SW101#cop nvram:startup-config tftp:
Address or name of remote host []? 10.1.0.1
Destination filename [sw101-confg]?
%Error opening tftp://10.1.0.1/sw101-confg (Permission denied)
SW101#

SW101#ping 10.1.0.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.0.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 3/7/15 ms
SW101#
```

Refer to the exhibit. SW101 could not transfer its startup configuration to a TFTP server. No ACL is configured on the switch, and it can successfully ping the host. Which action resolves the issue?

A. Open UDP port 69 on the TFTP server.

B. Open UDP port 179 on the TFTP server.

C. Configure a FW in the middle to allow bidirectional communication for TFTP.

D. Start the TFTP server on the host.

Correct Answer: *D*

DeWalt95 2 days, 14 hours ago

Selected Answer: D

Port being blocked would return a socket error?

upvoted 1 times

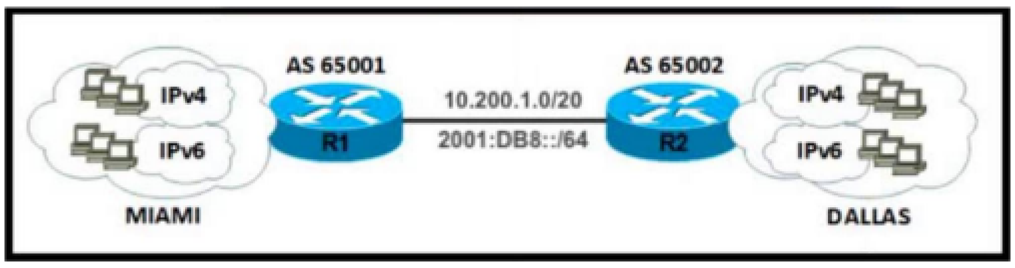

Refer to the exhibit. A network engineer configured routers R1 and R2 with MP-BGP. The engineer noticed that the routers cannot exchange any IPv6 routes; however, the IPv4 neighbor relationship is working fine. Which configuration must the engineer apply to router R2 to exchange IPv6 routes?

A.

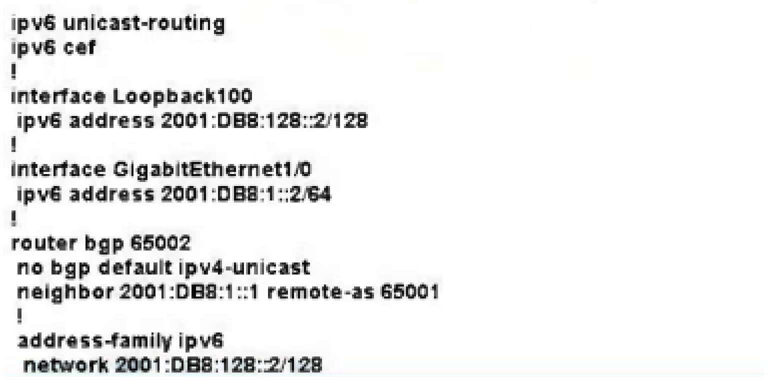

```
ipv6 unicast-routing
ipv6 cef
!
interface Loopback100
 ipv6 address 2001:DB8:128::2/128
!
interface GigabitEthernet1/0
 ipv6 address 2001:DB8:1::2/64
!
router bgp 65002
 no bgp default ipv4-unicast
 neighbor 2001:DB8:1::1 remote-as 65001
 !
 address-family ipv6
  network 2001:DB8:128::2/128
```

B.

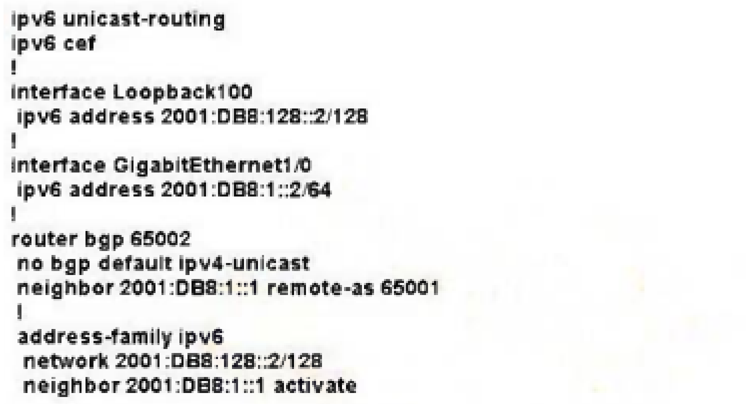

```
ipv6 unicast-routing
ipv6 cef
!
interface Loopback100
 ipv6 address 2001:DB8:128::2/128
!
interface GigabitEthernet1/0
 ipv6 address 2001:DB8:1::2/64
!
router bgp 65002
 no bgp default ipv4-unicast
 neighbor 2001:DB8:1::1 remote-as 65001
 !
 address-family ipv6
  network 2001:DB8:128::2/128
  neighbor 2001:DB8:1::1 activate
```

C.

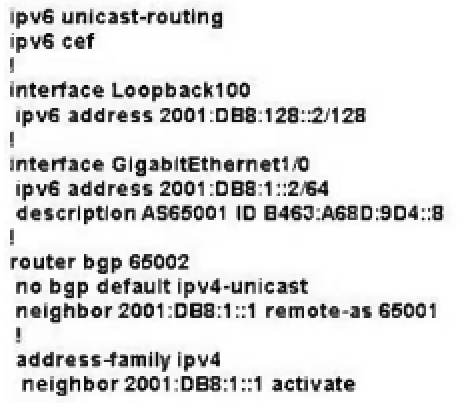

```
ipv6 unicast-routing
ipv6 cef
!
interface Loopback100
 ipv6 address 2001:DB8:128::2/128
!
interface GigabitEthernet1/0
 ipv6 address 2001:DB8:1::2/64
 description AS65001 ID B463:A68D:9D4::8
!
router bgp 65002
 no bgp default ipv4-unicast
 neighbor 2001:DB8:1::1 remote-as 65001
 !
 address-family ipv4
  neighbor 2001:DB8:1::1 activate
```

D.

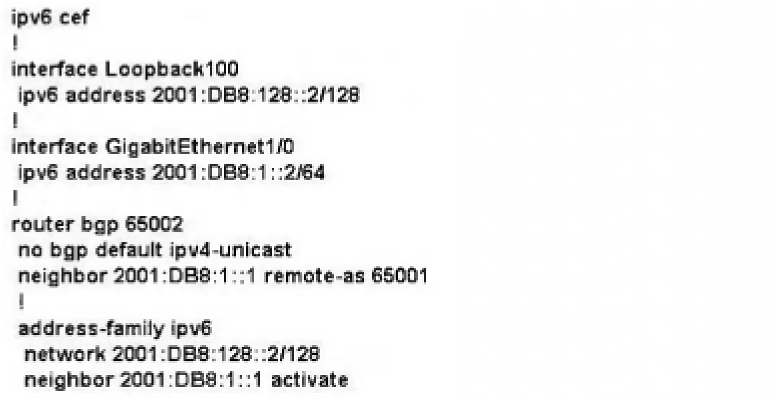

```
ipv6 cef
!
interface Loopback100
 ipv6 address 2001:DB8:128::2/128
!
interface GigabitEthernet1/0
 ipv6 address 2001:DB8:1::2/64
!
router bgp 65002
 no bgp default ipv4-unicast
 neighbor 2001:DB8:1::1 remote-as 65001
 !
 address-family ipv6
  network 2001:DB8:128::2/128
  neighbor 2001:DB8:1::1 activate
```

Correct Answer: *B*

DeWalt95 2 weeks, 3 days ago

Selected Answer: B

B is the only correct config.

upvoted 2 times

Question #513 *Topic 1*

What is an advantage of MPLS Layer 3 VPN deployment?

A. Planning and modifications are required for the customer intranet before migrating to Layer 3 VPN.

B. Scalable VPNs are created using connection-oriented, point-to-point, or multipoint overlay connections.

C. QoS provides performance with policy and support for a best-effort service level in an MPLS VPN.

D. Security is provided at the edge of the provider network through encryption.

Correct Answer: *A*

cebra 2 weeks, 3 days ago

Selected Answer: B

best effort is like no qos i stay with B

upvoted 1 times

> **cebra** 2 weeks, 3 days ago
>
> after more thinking the word connection-oriented is disturbing me
>
> IPSEC could be integrated between PE-PE and CE-CE Routers
> wich leads to answer D
> https://www.firewall.cx/cisco/cisco-routers/mpls-ip-vpn-security.html
> https://community.cisco.com/t5/service-providers-knowledge-base/static-ipsec-with-mpls-vpn/ta-p/3161428
>
> upvoted 1 times

DeWalt95 2 weeks, 3 days ago

Selected Answer: B

Would probably go for B but arguably C is also right

upvoted 1 times

> **DeWalt95** 3 days, 15 hours ago
>
> Upon thought - B seems to describe DMVPN. Change answer to C as traffic engineering is a part of MPLS.
>
> upvoted 1 times

RouterToRooter 2 weeks, 4 days ago

B or C

upvoted 1 times

Bombbear_W 4 weeks ago

The answer is C

upvoted 1 times

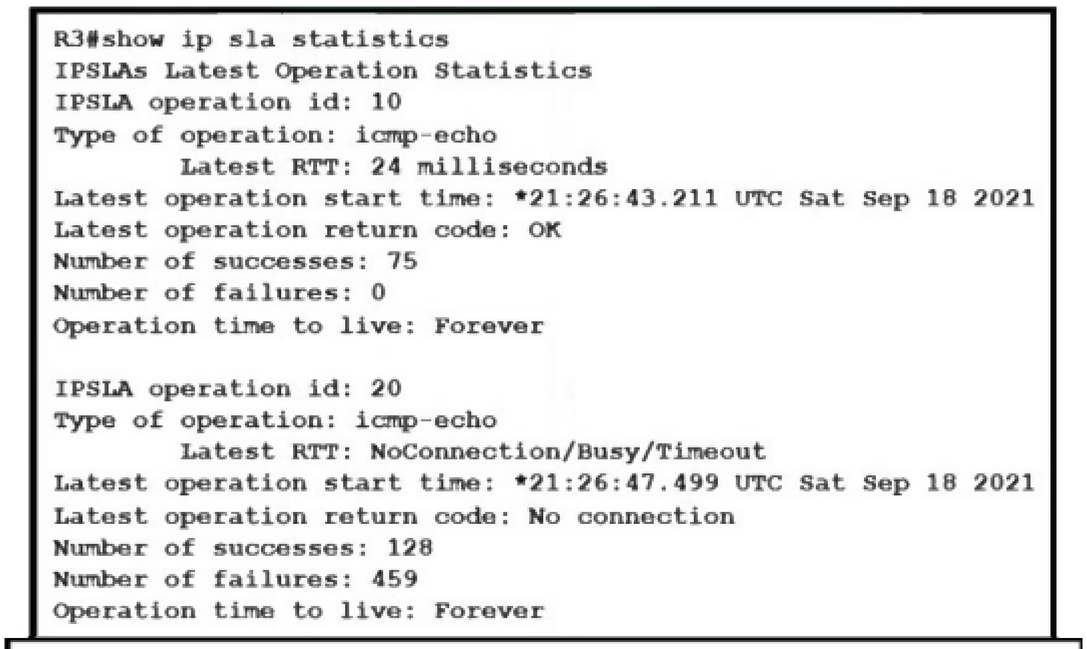

```
R3#show ip sla statistics
IPSLAs Latest Operation Statistics
IPSLA operation id: 10
Type of operation: icmp-echo
        Latest RTT: 24 milliseconds
Latest operation start time: *21:26:43.211 UTC Sat Sep 18 2021
Latest operation return code: OK
Number of successes: 75
Number of failures: 0
Operation time to live: Forever

IPSLA operation id: 20
Type of operation: icmp-echo
        Latest RTT: NoConnection/Busy/Timeout
Latest operation start time: *21:26:47.499 UTC Sat Sep 18 2021
Latest operation return code: No connection
Number of successes: 128
Number of failures: 459
Operation time to live: Forever
```

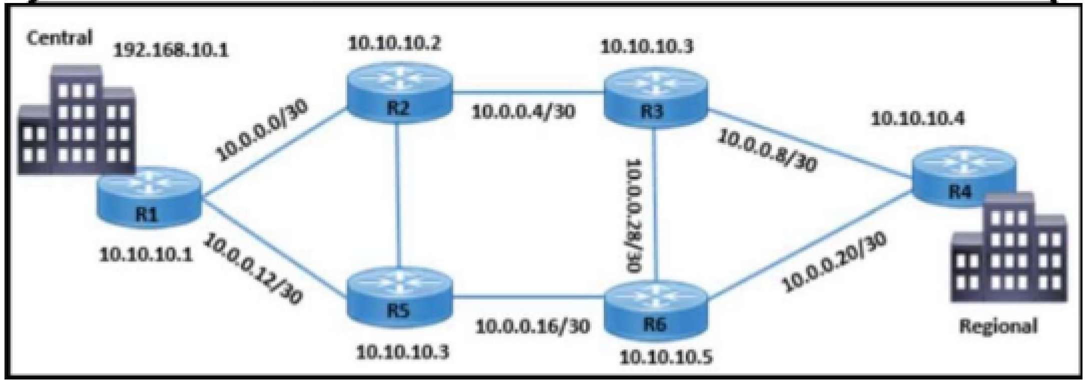

Refer to the exhibit. Traffic from R3 to the central site does not use alternate paths when R3 cannot reach 10.10.10.2. Traffic on R3 destined to R4 takes an alternate route via 10.10.10.6 when 10.10.10.4 is not accessible from R3. Which configuration switches traffic destined to 10.10.10.2 from R3 on the alternate path?

A. R3(config)#ip route 192.168.10.1 255.255.255.255 10.10.10.2 track 20

B. R6(config)#ip route 10.10.10.3 255.255.255.255 10.0.0.30

C. R3(config)#track 20 ip sla 20 reachability

D. R2(config)#ip route 10.10.10.3 255.255.255.255 10.0.0.6

Correct Answer: *A*

DeWalt95 2 weeks, 3 days ago

Selected Answer: A

Perhaps missing information? But A is the only command that makes any sense even though we cant see the SLA targets.

upvoted 2 times

Question #515 Topic 1

A newly installed router starts establishing an LDP session from another MPLS router to which it is not directly connected. Which LDP message type responds by target router to the initiating router using UDP protocol?

A. notification message

B. session message

C. advertisement message

D. extended discovery message

Correct Answer: *D*

DeWalt95 2 weeks, 3 days ago

Selected Answer: C

Agree its D based on the link Tedmus provided.
Think it unreasonable to ask questions about specific MPLS configurations in a Enterprise exam.

upvoted 1 times

av3672 3 weeks, 4 days ago

The correct answer is:

B. session message

When a newly installed router initiates the establishment of an LDP (Label Distribution Protocol) session to another MPLS router, the target router responds with a session message. This message type is part of the LDP session establishment process and is exchanged between the routers to establish the necessary parameters for label distribution.

Therefore, the correct option is B. session message.

upvoted 1 times

Tedmus 3 weeks, 5 days ago

Selected Answer: D

Correct anwer is D.
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/convert/mp_ldp_book/mp_ldp_overview.html

Nondirectly Connected MPLS LDP Sessions

If the LSR is more than one hop from its neighbor, it is nondirectly connected to its neighbor. For these nondirectly connected neighbors, the LSR sends out a targeted Hello message as a UDP packet, but as a unicast message specifically addressed to that LSR. The nondirectly connected LSR responds to the Hello message and the two routers begin to establish an LDP session. This is called extended discovery.

upvoted 1 times

Bombbear_W 4 weeks ago

The correct answer is:
C. advertisement message

In MPLS (Multiprotocol Label Switching), LDP (Label Distribution Protocol) is used to establish label-switched paths (LSPs) and distribute labels between routers. When a newly installed router starts establishing an LDP session with another MPLS router to which it is not directly connected, the target router responds with an LDP advertisement message. This message is part of the label distribution process and is used to exchange information about labels and establish the LSP.
So, the correct option is C. advertisement message.

upvoted 4 times

Question #516 *Topic 1*

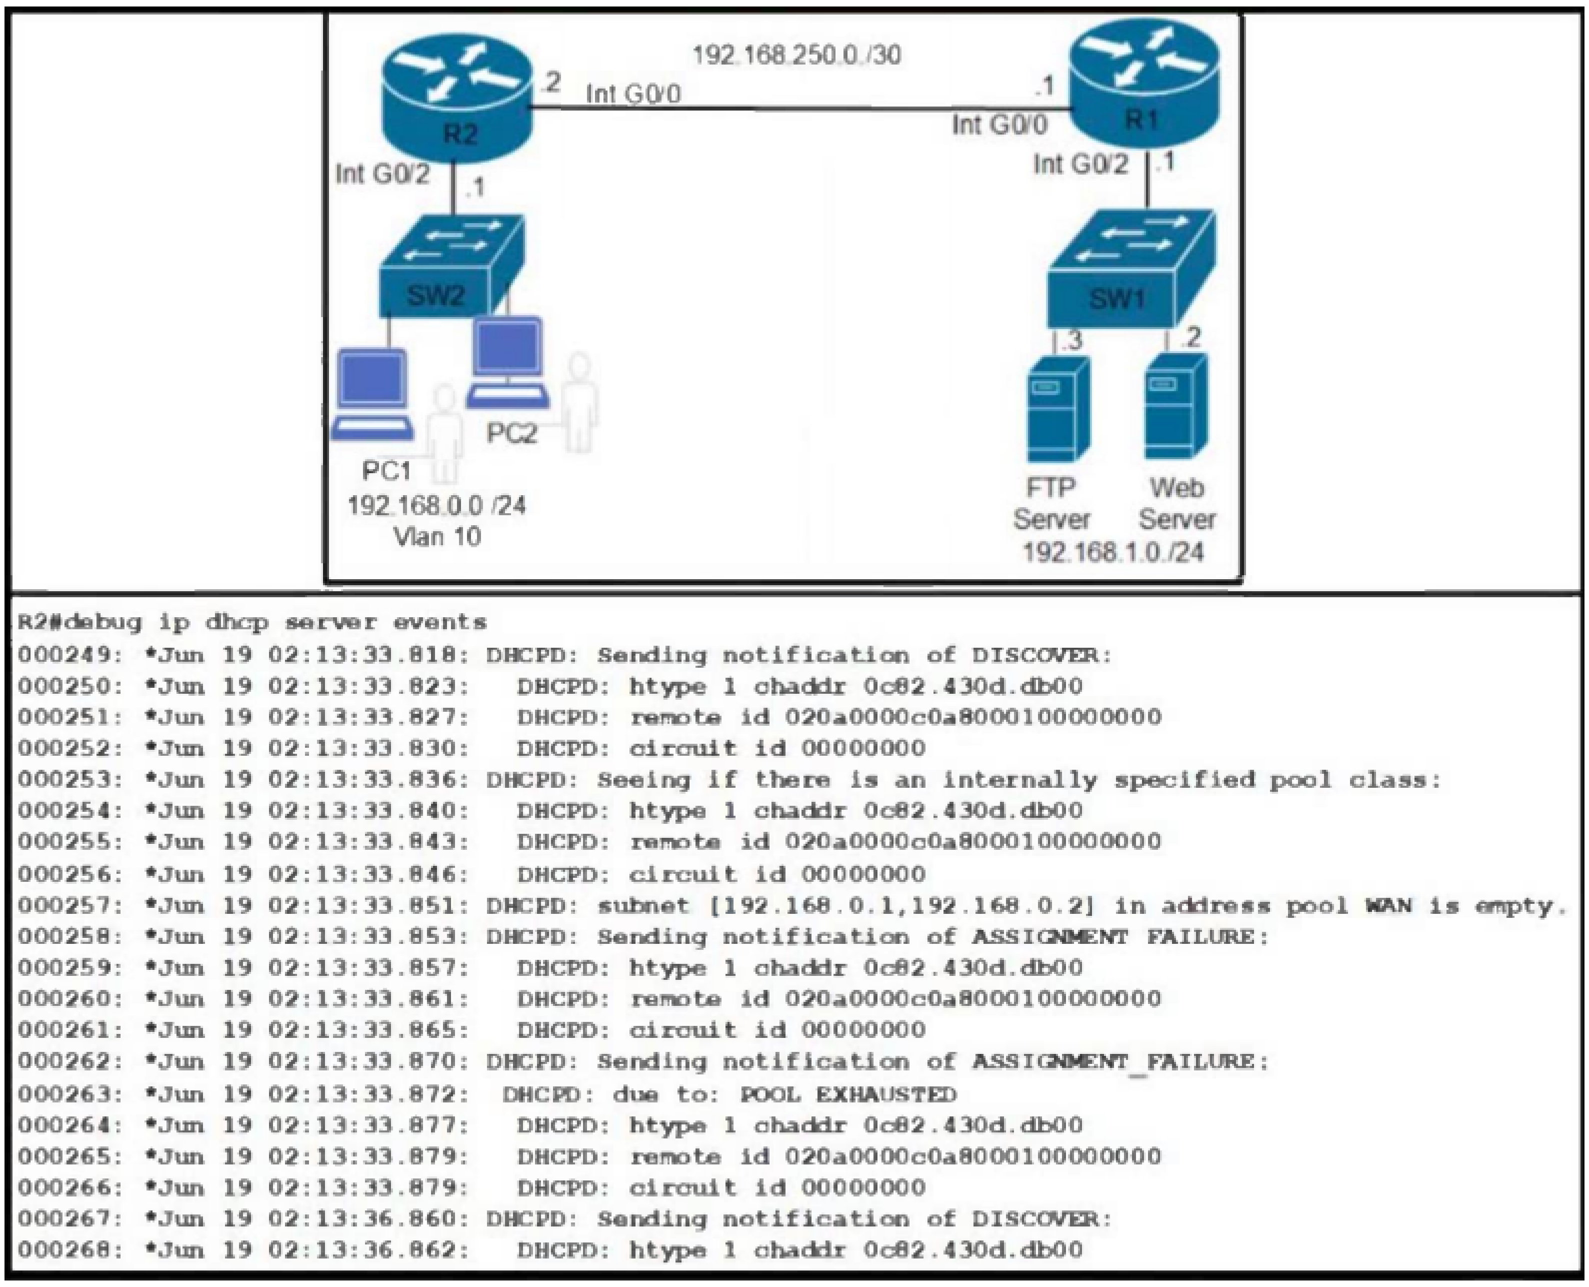

```
R2#debug ip dhcp server events
000249: *Jun 19 02:13:33.818: DHCPD: Sending notification of DISCOVER:
000250: *Jun 19 02:13:33.823:   DHCPD: htype 1 chaddr 0c82.430d.db00
000251: *Jun 19 02:13:33.827:   DHCPD: remote id 020a0000c0a800010000000
000252: *Jun 19 02:13:33.830:   DHCPD: circuit id 00000000
000253: *Jun 19 02:13:33.836: DHCPD: Seeing if there is an internally specified pool class:
000254: *Jun 19 02:13:33.840:   DHCPD: htype 1 chaddr 0c82.430d.db00
000255: *Jun 19 02:13:33.843:   DHCPD: remote id 020a0000c0a800010000000
000256: *Jun 19 02:13:33.846:   DHCPD: circuit id 00000000
000257: *Jun 19 02:13:33.851: DHCPD: subnet [192.168.0.1,192.168.0.2] in address pool WAN is empty.
000258: *Jun 19 02:13:33.853: DHCPD: Sending notification of ASSIGNMENT FAILURE:
000259: *Jun 19 02:13:33.857:   DHCPD: htype 1 chaddr 0c82.430d.db00
000260: *Jun 19 02:13:33.861:   DHCPD: remote id 020a0000c0a800010000000
000261: *Jun 19 02:13:33.865:   DHCPD: circuit id 00000000
000262: *Jun 19 02:13:33.870: DHCPD: Sending notification of ASSIGNMENT_FAILURE:
000263: *Jun 19 02:13:33.872:  DHCPD: due to: POOL EXHAUSTED
000264: *Jun 19 02:13:33.877:   DHCPD: htype 1 chaddr 0c82.430d.db00
000265: *Jun 19 02:13:33.879:   DHCPD: remote id 020a0000c0a800010000000
000266: *Jun 19 02:13:33.879:   DHCPD: circuit id 00000000
000267: *Jun 19 02:13:36.860: DHCPD: Sending notification of DISCOVER:
000268: *Jun 19 02:13:36.862:   DHCPD: htype 1 chaddr 0c82.430d.db00
```

Refer to the exhibit. Router R2 VLAN 10 users cannot get dynamic IP addresses from R1. Which action resolves the issue?

A. Eliminate the port security feature on the ports of switch SW2.

B. Identify the host with the duplicate IP address.

C. Configure the IP helper feature on the Interface GigabitEthernet 0/2 of router R2.

D. Expand the address scope of VLAN 10.

Correct Answer: *D*

DeWalt95 3 days, 15 hours ago

Selected Answer: B

Changing answer to B. Its a error related to duplicate addresses

upvoted 1 times

Tedmus 1 week, 6 days ago

Selected Answer: B

It's a trap.
Not D, because we have 192.168.0.1 - 192.168.0.2 to assign and only two users.
So ip-helper is working.
Found several entries at Cisco to explain that duplicate IP addresses causes this error.

upvoted 1 times

DeWalt95 3 days, 15 hours ago

Excellent spot

upvoted 1 times

DeWalt95 2 weeks, 3 days ago

Selected Answer: D

Error messages pretty explicit - none of the other answers make sense.

upvoted 1 times

DeWalt95 3 days, 15 hours ago

Ignore this - I was wrong
upvoted 1 times

Question #517 *Topic 1*

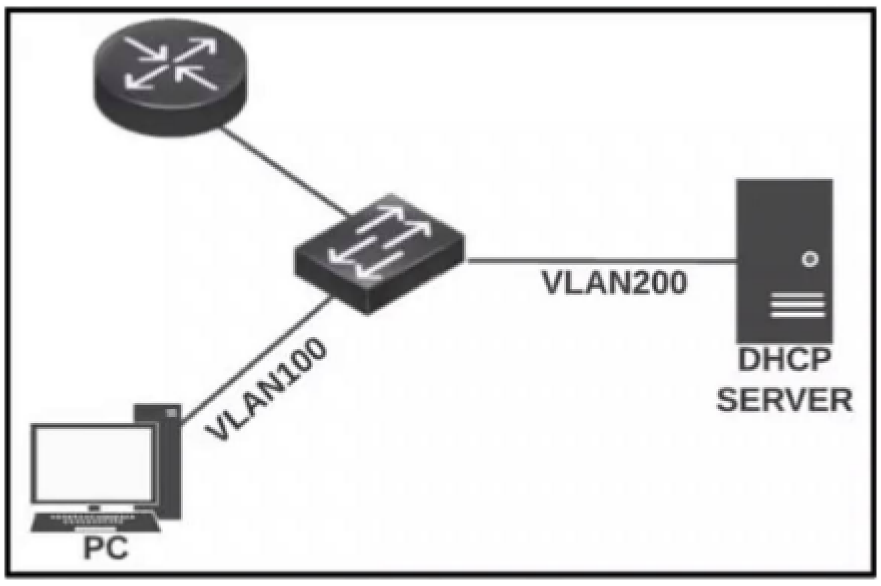

Refer to the exhibit. A PC is configured to obtain an IP address automatically, but it receives an IP address only from the 169.254.0.0 subnet. The DHCP server logs contained no DHCPDISCOVER message from the MAC address of the PC. Which action resolves the issue?

A. Configure a DHCP reservation on the server for the PC.

B. Configure an ip helper-address on the router to forward DHCP messages to the server.

C. Configure DHCP Snooping on the switch to forward DHCP messages to the server.

D. Configure a static IP address on the PC and exclude it from the DHCP pool.

Correct Answer: *B*

DeWalt95 2 weeks, 3 days ago

Selected Answer: B

DHCP helper command needed as Client on a different VLAN to the DHCP server
upvoted 1 times

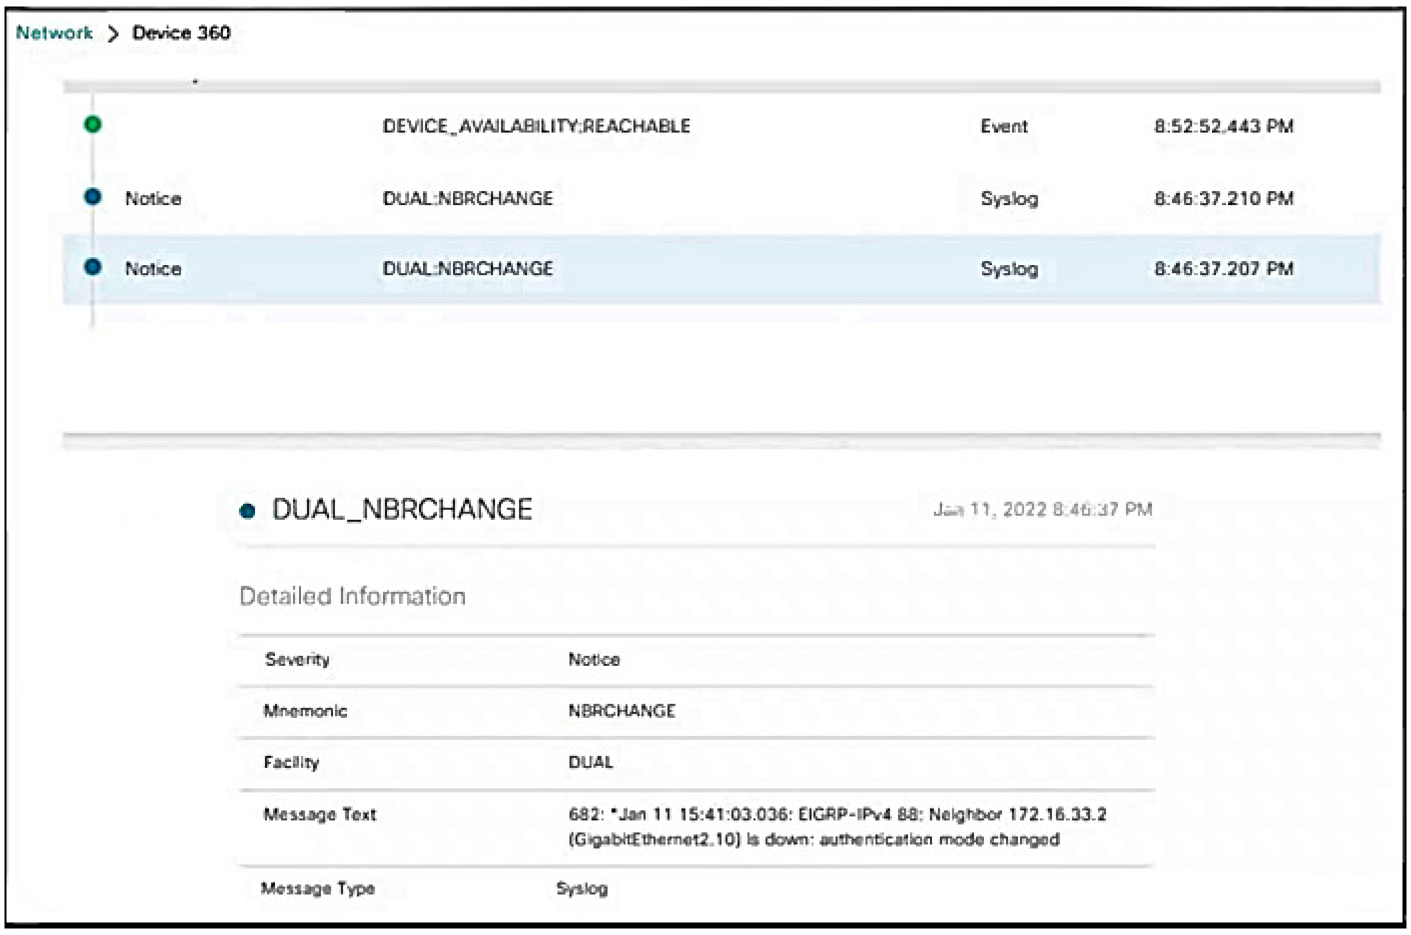

Network > Device 360

●	DEVICE_AVAILABILITY:REACHABLE	Event	8:52:52.443 PM
● Notice	DUAL:NBRCHANGE	Syslog	8:46:37.210 PM
● Notice	DUAL:NBRCHANGE	Syslog	8:46:37.207 PM

● DUAL_NBRCHANGE Jan 11, 2022 8:46:37 PM

Detailed Information

Severity	Notice
Mnemonic	NBRCHANGE
Facility	DUAL
Message Text	682: *Jan 11 15:41:03.036: EIGRP-IPv4 88: Neighbor 172.16.33.2 (GigabitEthernet2.10) is down: authentication mode changed
Message Type	Syslog

Refer to the exhibit. R1 lost its directly connected EIGRP peer 172.16.33.2 (SW1). Which configuration resolves the issue?

A.

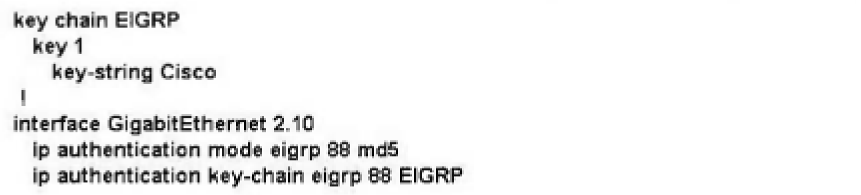

```
key chain EIGRP
  key 1
    key-string Cisco
 !
interface GigabitEthernet 2.10
  ip authentication mode eigrp 88 md5
  ip authentication key-chain eigrp 88 EIGRP
```

B.

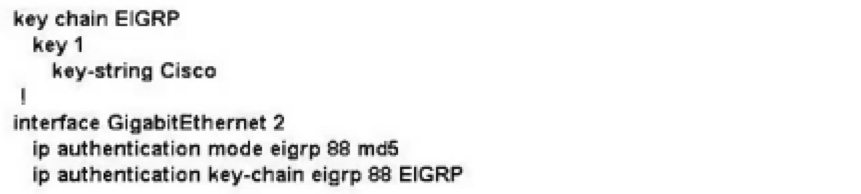

```
key chain EIGRP
  key 1
    key-string Cisco
 !
interface GigabitEthernet 2
  ip authentication mode eigrp 88 md5
  ip authentication key-chain eigrp 88 EIGRP
```

C.

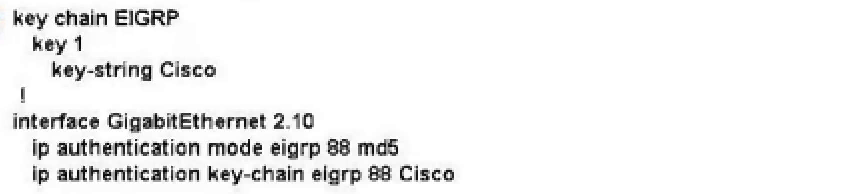

```
key chain EIGRP
  key 1
    key-string Cisco
 !
interface GigabitEthernet 2.10
  ip authentication mode eigrp 88 md5
  ip authentication key-chain eigrp 88 Cisco
```

D.

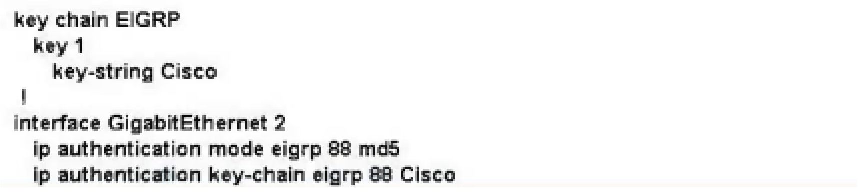

```
key chain EIGRP
  key 1
    key-string Cisco
 !
interface GigabitEthernet 2
  ip authentication mode eigrp 88 md5
  ip authentication key-chain eigrp 88 Cisco
```

Correct Answer: *C*

⊟ **DeWalt95** 2 weeks, 3 days ago

Selected Answer: A

Its a basic spot the difference challenge and the answer is A

upvoted 2 times

⊟ **sayed_2908** 4 weeks, 1 day ago

Selected Answer: A

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/82110-eigrp-authentication.html

Answer A

upvoted 3 times

⊟ **b8os5h** 4 weeks, 1 day ago

Selected Answer: A

(config-if)#ip authentication key-chain eigrp < AS > < key-chain-name >

upvoted 3 times

Question #519

Topic 1

How are CE advertised routes segmented from other CE routers on an MPLS PE router?

A. with a combination of VRF-Lite and MP-BGP

B. by pushing MPLS labels advertised by LDP on customer routes

C. by enabling multiple instances of BGP, one for each CE router

D. by assigning CE-facing interfaces to different VRFs

Correct Answer: *A*

DeWalt95 2 weeks, 3 days ago

Selected Answer: D

Agree answer is D but a poorly worded question

upvoted 1 times

RouterToRooter 2 weeks, 4 days ago

D is the answer

upvoted 1 times

av3672 3 weeks, 4 days ago

Selected Answer: D

D. by assigning CE-facing interfaces to different VRFs

upvoted 3 times

Bombbear_W 4 weeks ago

D. by assigning CE-facing interfaces to different VRFs

In an MPLS (Multiprotocol Label Switching) network, Customer Edge (CE) routers are typically connected to Provider Edge (PE) routers. To keep the routes from different CE routers segregated on the PE router, Virtual Routing and Forwarding (VRF) is commonly used. Each VRF acts as a separate routing table, effectively isolating routes from one CE router from routes of another CE router.

So, the correct answer is D.

upvoted 1 times

b8os5h 4 weeks, 1 day ago

Selected Answer: D

The question says "on an MPLS PE router", so the question asks how to segment inside a single router.

upvoted 1 times

Question #520

Topic 1

Which Layer 3 VPN attribute installs customer routes in the VRF?

A. RD

B. RT

C. extended-community

D. MPLS label

Correct Answer: *B*

AlexInShort12 4 days, 6 hours ago

L3 VPN Attribute:
RD (Route Distinguisher)
IPv4 prefix
Next Hop
VPN Label
Bad question, extended-community contain the RT.

upvoted 1 times

DeWalt95 2 weeks, 3 days ago

Selected Answer: B

Given answer is correct

upvoted 2 times

Question #521 Topic 1

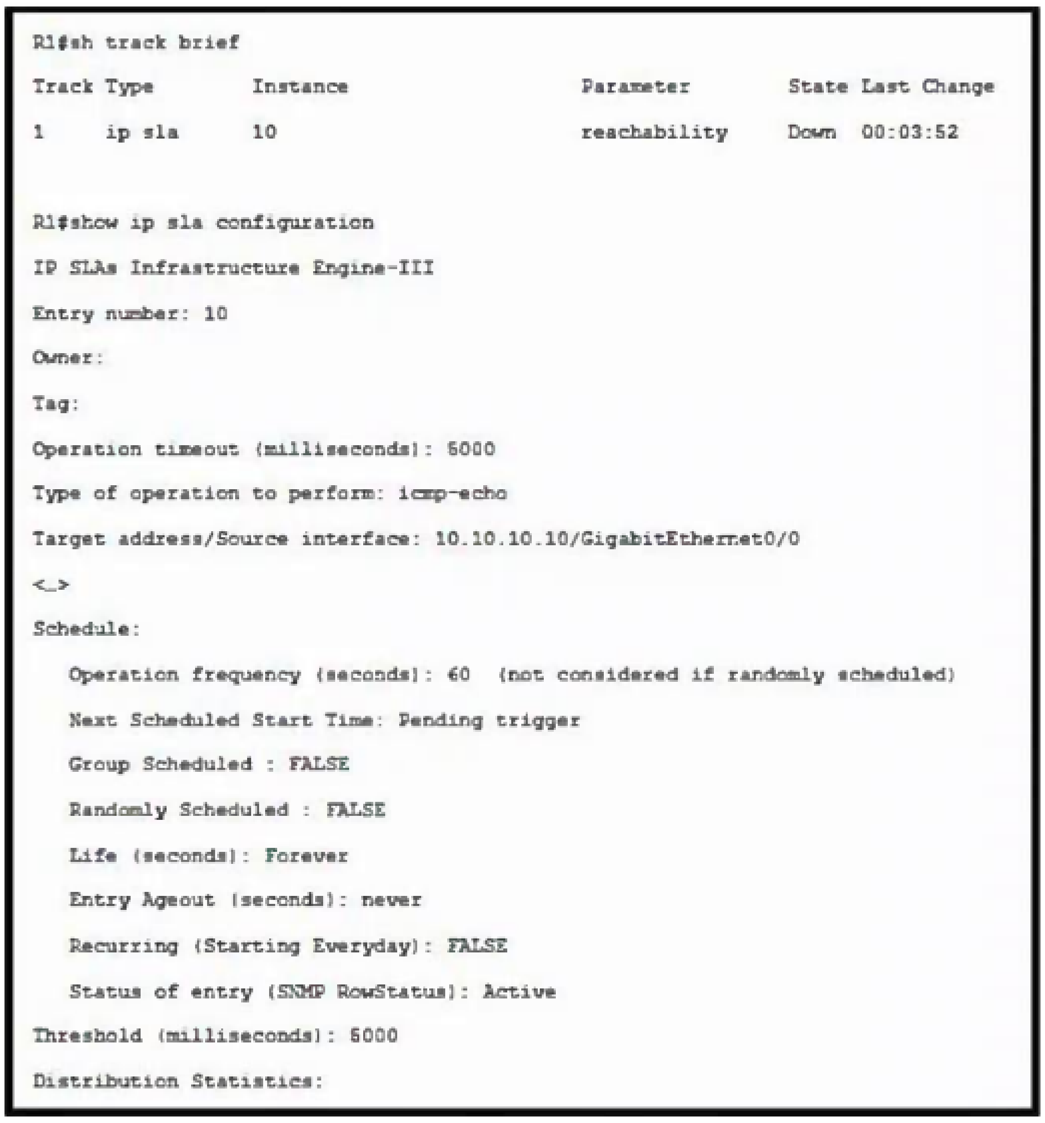

```
R1#sh track brief
Track Type       Instance                   Parameter       State Last Change
1     ip sla     10                         reachability    Down  00:03:52


R1#show ip sla configuration
IP SLAs Infrastructure Engine-III
Entry number: 10
Owner:
Tag:
Operation timeout (milliseconds): 5000
Type of operation to perform: icmp-echo
Target address/Source interface: 10.10.10.10/GigabitEthernet0/0
<_>
Schedule:
   Operation frequency (seconds): 60  (not considered if randomly scheduled)
   Next Scheduled Start Time: Pending trigger
   Group Scheduled : FALSE
   Randomly Scheduled : FALSE
   Life (seconds): Forever
   Entry Ageout (seconds): never
   Recurring (Starting Everyday): FALSE
   Status of entry (SNMP RowStatus): Active
Threshold (milliseconds): 5000
Distribution Statistics:
```

Refer to the exhibit. A network engineer notices that the configured track option is down. Which configuration resolves the issue?

A. ip sla schedule 10 start-time pending life forever

B. ip sla schedule 10
no timeout

C. ip sla schedule 10 start-time now

D. ip sla schedule 10
no threshold

Correct Answer: *C*

Question #522

Topic 1

Which technique removes the outermost label of an MPLS-tagged packet before the packet is forwarded to an adjacent LER?

A. explicit-null

B. PHP

C. label swap

D. label imposition

Correct Answer: *B*

⊟ **RouterToRooter** 2 weeks, 4 days ago

Selected Answer: B

Penultimate hop popping (PHP) is a function performed by certain routers in an MPLS enabled network. It refers to the process whereby the outermost label of an MPLS tagged packet is removed by a label switch router (LSR) before the packet is passed to an adjacent label edge router (LER).

upvoted 1 times

Question #523 Topic 1

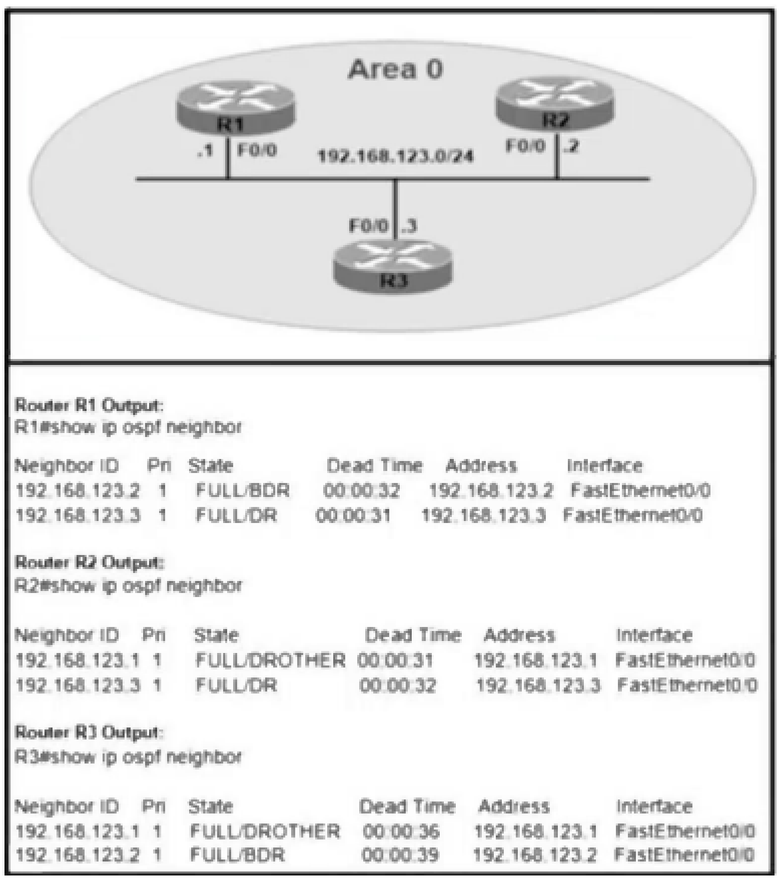

```
Router R1 Output:
R1#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
192.168.123.2    1    FULL/BDR        00:00:32    192.168.123.2   FastEthernet0/0
192.168.123.3    1    FULL/DR         00:00:31    192.168.123.3   FastEthernet0/0

Router R2 Output:
R2#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
192.168.123.1    1    FULL/DROTHER    00:00:31    192.168.123.1   FastEthernet0/0
192.168.123.3    1    FULL/DR         00:00:32    192.168.123.3   FastEthernet0/0

Router R3 Output:
R3#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
192.168.123.1    1    FULL/DROTHER    00:00:36    192.168.123.1   FastEthernet0/0
192.168.123.2    1    FULL/BDR        00:00:39    192.168.123.2   FastEthernet0/0
```

Refer to the exhibit. An administrator wanted to make R1 always elected as DR, R2 as BDR, and R3 as DROTHER but could not achieve the desired results. Which two configurations resolve the issue? (Choose two.)

A. On the R3 F0/0 interface, configure OSPFpriority to 201.

B. On the R1 F0/0 interface, configure OSPFpriority to 202.

C. On the R2 F0/0 interface, configure OSPFpriority to 200.

D. On the R1 F0/0 interface, configure OSPFpriority to 255.

E. On the R2 F0/0 interface, configure OSPFpriority to 201.

Correct Answer: *BE*

⊟ **DeWalt95** 2 weeks, 2 days ago

Selected Answer: DE

Agree its DE..the make sure R1 is 'Always' the DR then use the highest possible priority.

upvoted 1 times

⊟ **Tedmus** 3 weeks, 3 days ago

Selected Answer: CD

To make R1 "ALWAYS" the DR set it better to 255.
Also set the R2 with priority of 200 very high.
Let R3 with defauklt priority of "1".
If you really not want to be elected set the priority to "0". But this is not an option with this question.

upvoted 2 times

⊟ **Tedmus** 3 weeks, 3 days ago

I mean D & E
Set R2 to 201.

upvoted 2 times

Question #524 Topic 1

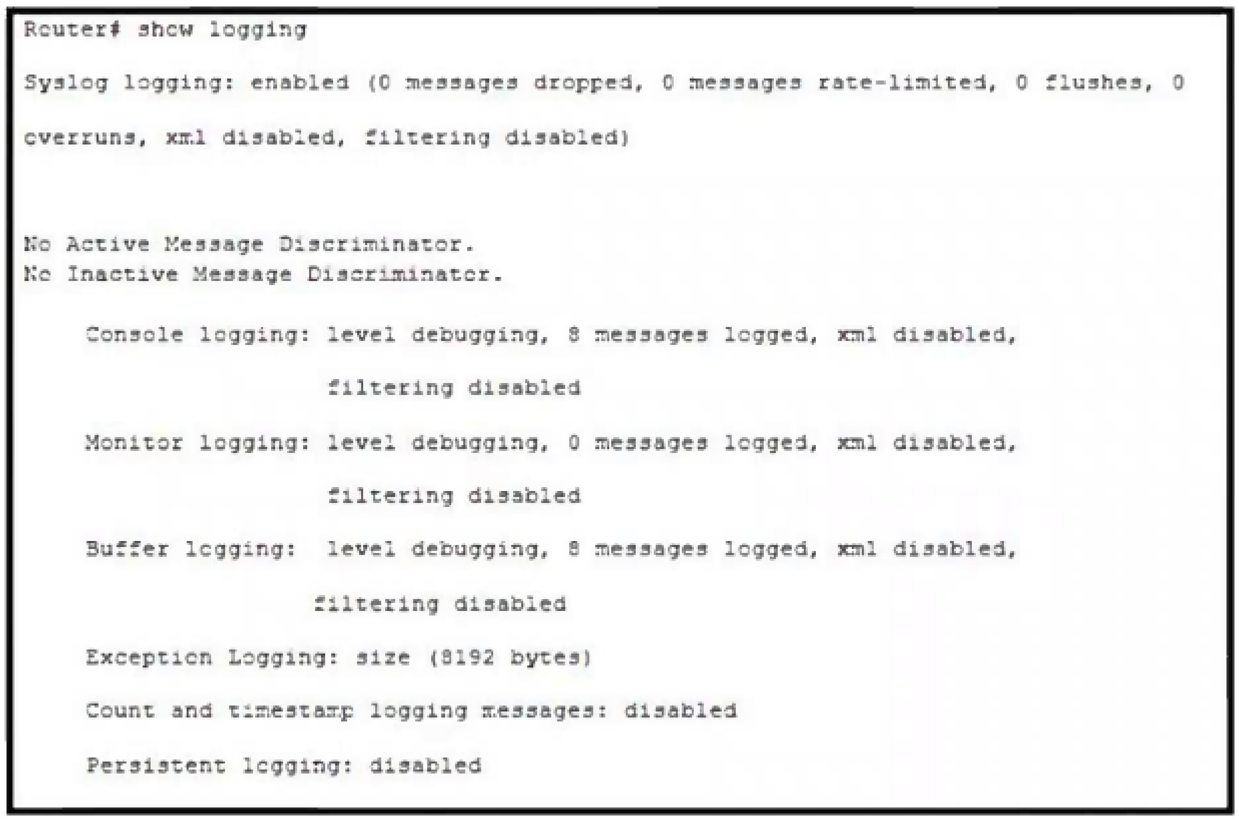

```
Router# show logging
Syslog logging: enabled (0 messages dropped, 0 messages rate-limited, 0 flushes, 0
overruns, xml disabled, filtering disabled)


No Active Message Discriminator.
No Inactive Message Discriminator.
    Console logging: level debugging, 8 messages logged, xml disabled,
                     filtering disabled
    Monitor logging: level debugging, 0 messages logged, xml disabled,
                     filtering disabled
    Buffer logging:  level debugging, 8 messages logged, xml disabled,
                    filtering disabled
    Exception Logging: size (8192 bytes)
    Count and timestamp logging messages: disabled
    Persistent logging: disabled
```

Refer to the exhibit. A network engineer lost remote access to the router due to a network problem. The engineer used the console to access the router and noticed continuous logs on the console terminal. Which configuration limits the number of log messages on the console to critical and higher severity level messages?

A. logging console 2

B. logging console 5

C. no logging console

D. term no monitor

Correct Answer: *A*

DeWalt95 2 weeks, 2 days ago

Selected Answer: A

A is correct

upvoted 1 times

RouterToRooter 2 weeks, 4 days ago

Selected Answer: A

0. Emergency
1. Alert
2. Critical
3. Error
4. Warning
5. Notice
6. Informational
7. Debug

upvoted 1 times

Question #525 *Topic 1*

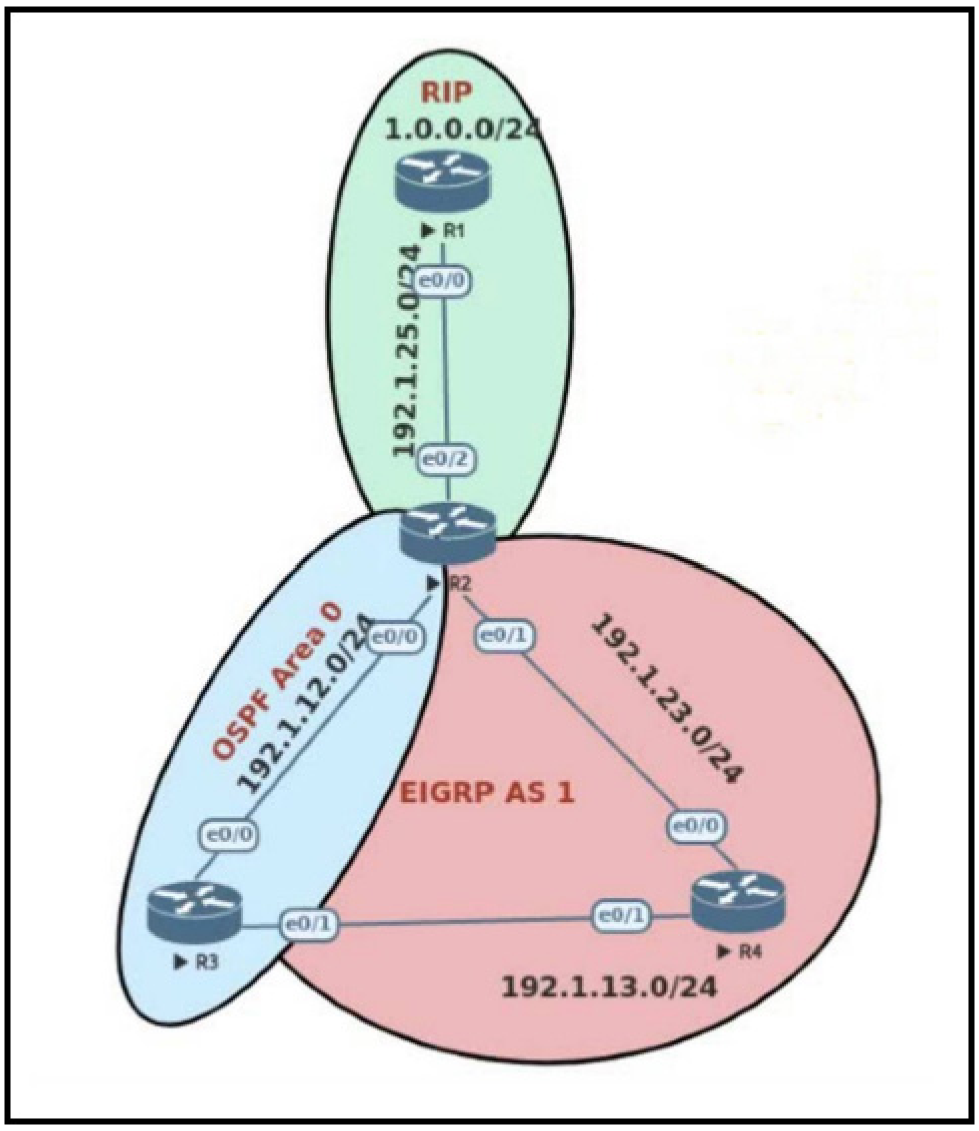

Refer to the exhibit. R3 is learning the 1.0.0.0/24 route through OSPF instead of EIGRP. Which action causes R3 to choose EIGRP to reach the 1.0.0.0/24 network?

A. Configure EIGRP administrative distance to 120.

B. Configure EIGRP administrative distance to 110.

C. Configure OSPF administrative distance to 120.

D. Configure OSPF administrative distance to 200.

Correct Answer: *D*

DeWalt95 2 days, 14 hours ago

Selected Answer: D

Must make OSPF AD more than EIGRP External 170

upvoted 1 times

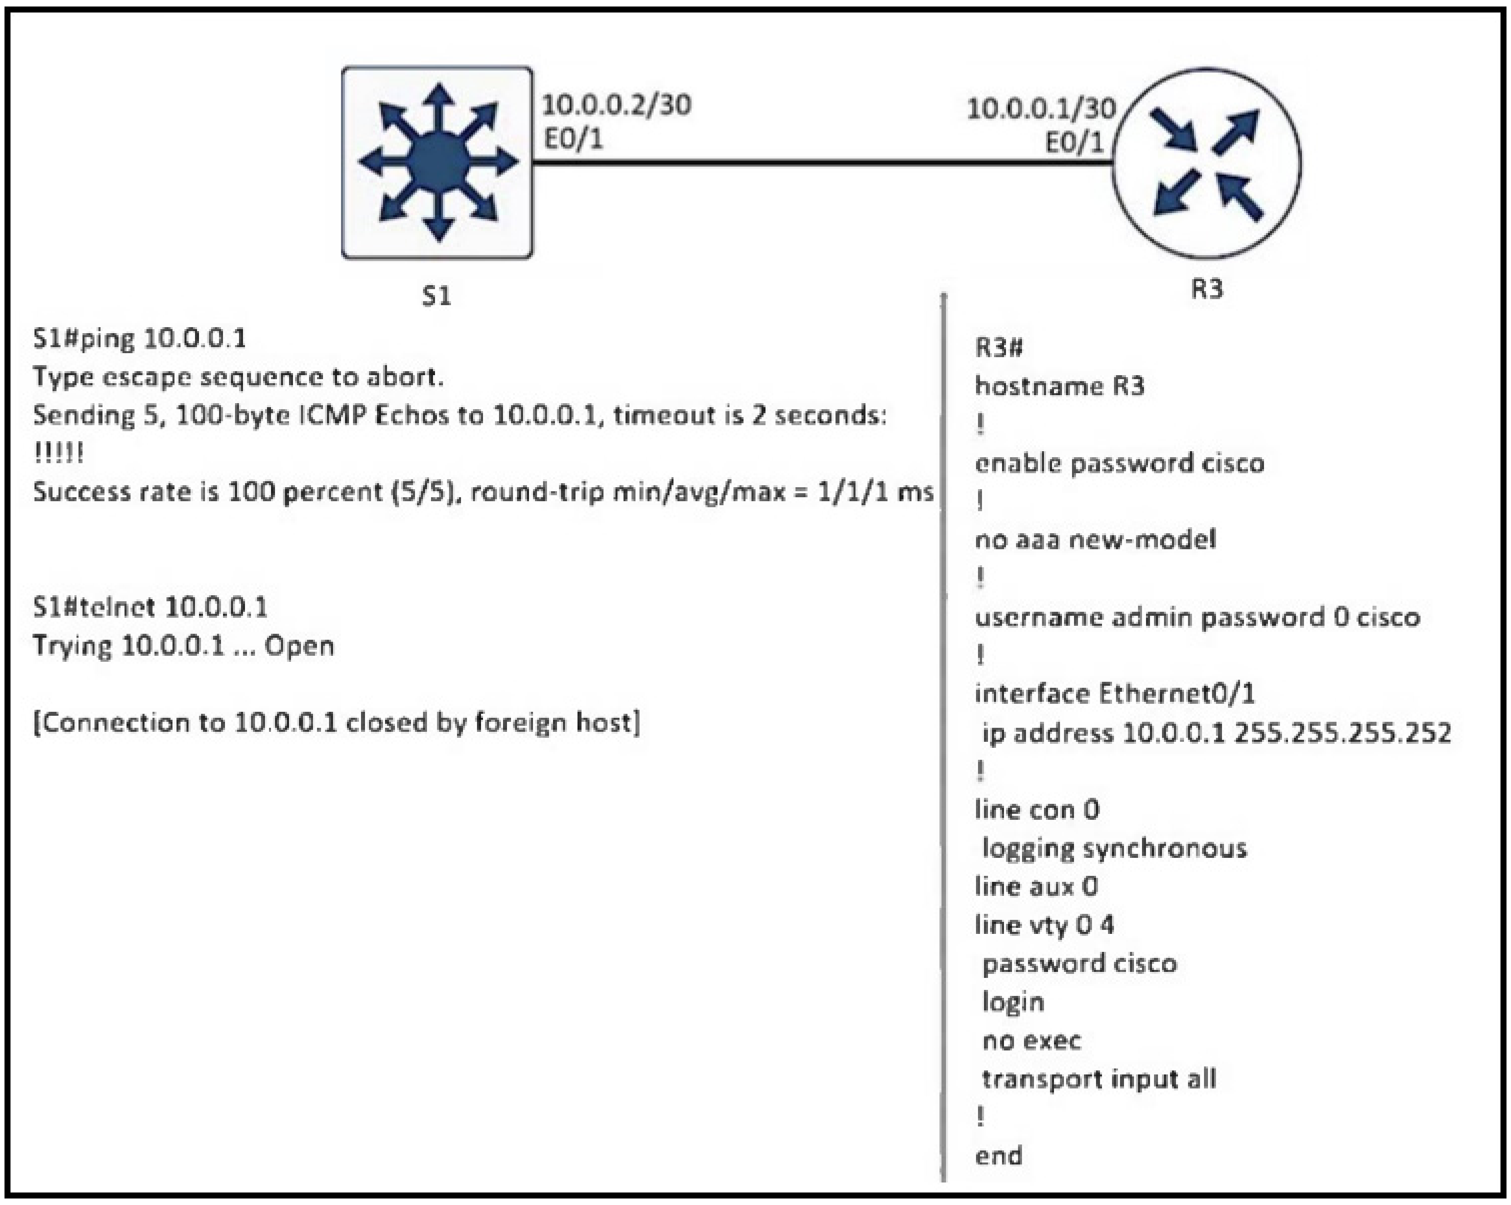

```
S1#ping 10.0.0.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.0.0.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms


S1#telnet 10.0.0.1
Trying 10.0.0.1 ... Open

[Connection to 10.0.0.1 closed by foreign host]
```

```
R3#
hostname R3
!
enable password cisco
!
no aaa new-model
!
username admin password 0 cisco
!
interface Ethernet0/1
 ip address 10.0.0.1 255.255.255.252
!
line con 0
 logging synchronous
line aux 0
line vty 0 4
 password cisco
 login
 no exec
 transport input all
!
end
```

Refer to the exhibit. A network engineer cannot remote access R3 using Telnet from switch S1. Which action resolves the issue?

A. Allow to use the ssh -| admin 10.0.0.1 command on the switch.

B. Add the login admin command on the switch.

C. Add the transport input telnet command on R3.

D. Allow the inbound connection via the exec command on R3.

Correct Answer: *D*

RouterToRooter 2 weeks, 4 days ago

Selected Answer: D

no Exec command on R3 stopping telnet

upvoted 1 times

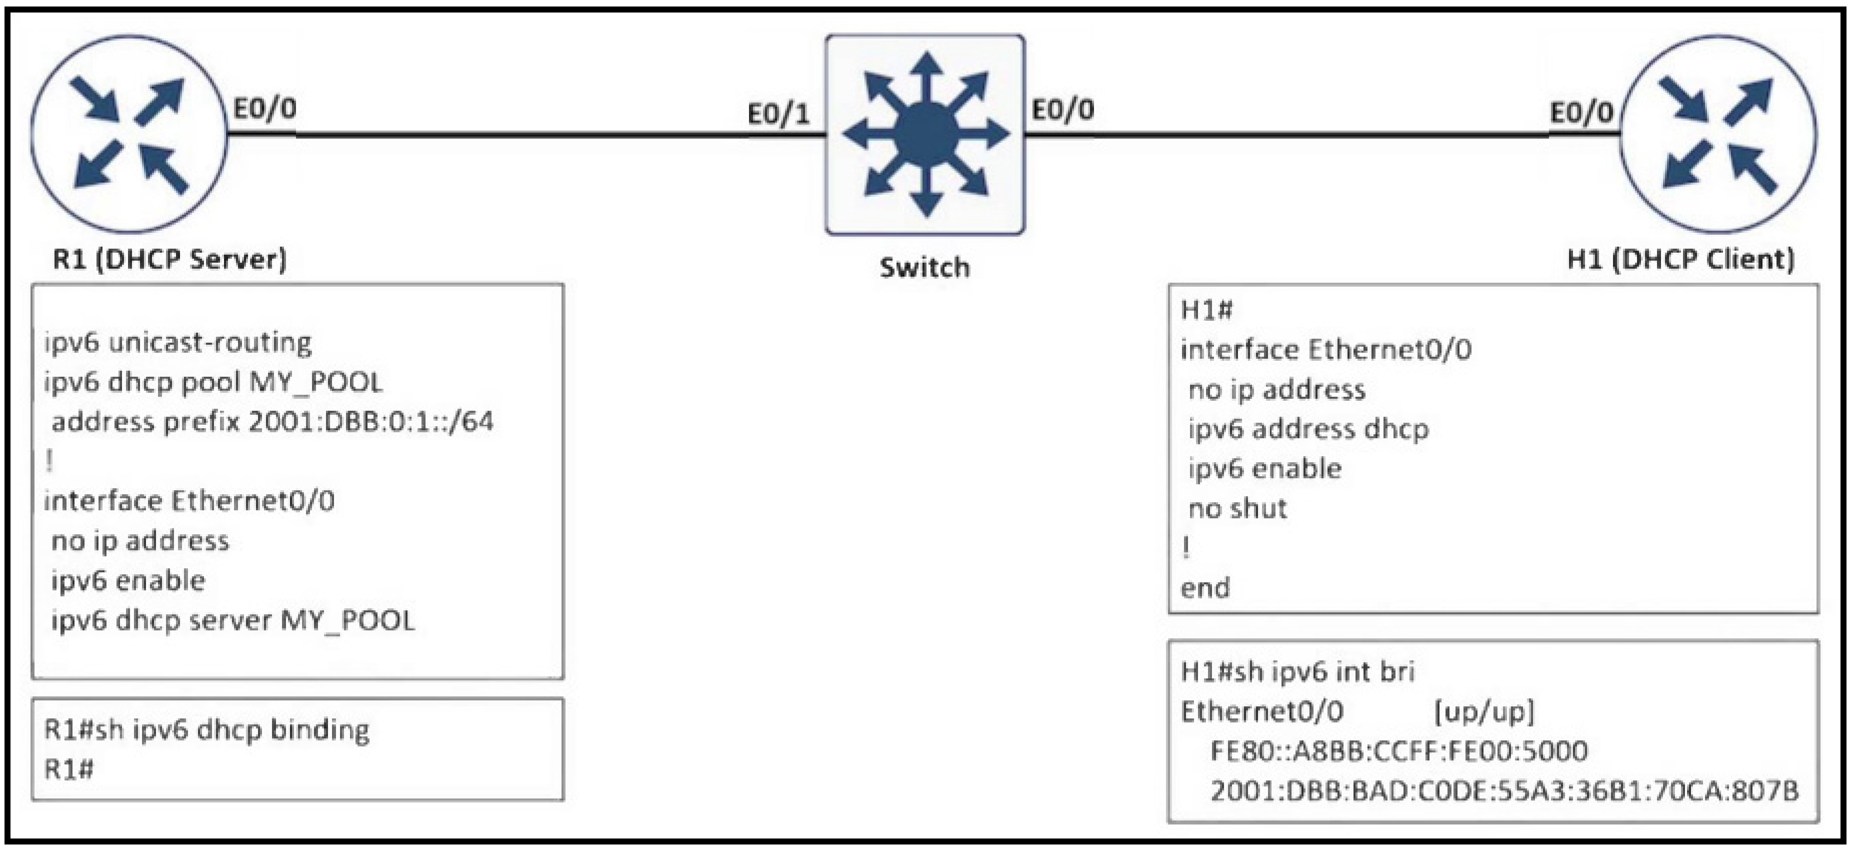

```
ipv6 unicast-routing
ipv6 dhcp pool MY_POOL
 address prefix 2001:DBB:0:1::/64
!
interface Ethernet0/0
 no ip address
 ipv6 enable
 ipv6 dhcp server MY_POOL
```

```
R1#sh ipv6 dhcp binding
R1#
```

```
H1#
interface Ethernet0/0
 no ip address
 ipv6 address dhcp
 ipv6 enable
 no shut
!
end
```

```
H1#sh ipv6 int bri
Ethernet0/0          [up/up]
   FE80::A8BB:CCFF:FE00:5000
   2001:DBB:BAD:CODE:55A3:36B1:70CA:807B
```

Refer to the exhibit. The client received the IPv6 address from the IPv6 DHCP server but the show command does not show the IPv6 DHCP bindings on the server. Which action resolves the issue?

A. Extend the DHCP lease time because R1 removed the IPv6 address earlier after the lease expired.

B. Configure H1 as the DHCP client that manually assigns the IPv6 address on interface e0/0.

C. Configure authorized DHCP servers to avoid IPv6 addresses from a rogue DHCP server.

D. Use the 2001:DBB:BAD:CODE::/64 prefix for the DHCP pool on R1.

Correct Answer: *D*

Tedmus Highly Voted 2 weeks, 5 days ago

Selected Answer: C

I guess we have a rogue DHCP server in the network - so "C" is more correct.
Topic is 3.4 IPv6 FHS.

upvoted 5 times

DeWalt95 Most Recent 2 days, 14 hours ago

Selected Answer: C

'Bad code' is a bit of a clue. Agree its an IPV6 FHP question not a DHCP one.

upvoted 1 times

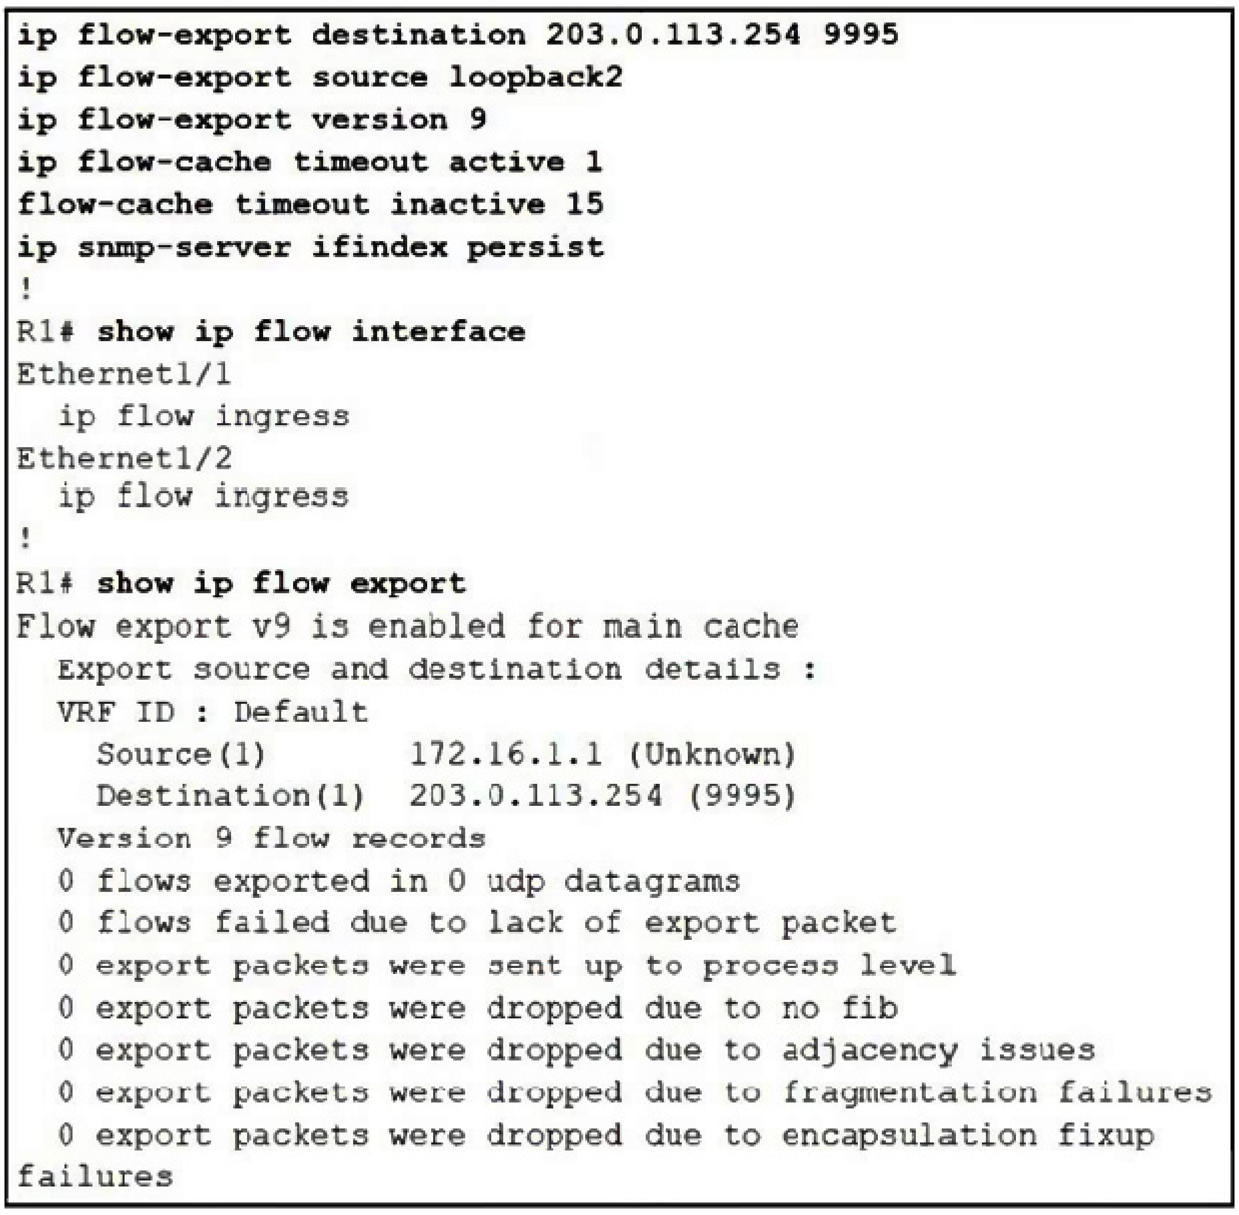

```
ip flow-export destination 203.0.113.254 9995
ip flow-export source loopback2
ip flow-export version 9
ip flow-cache timeout active 1
flow-cache timeout inactive 15
ip snmp-server ifindex persist
!
R1# show ip flow interface
Ethernet1/1
  ip flow ingress
Ethernet1/2
  ip flow ingress
!
R1# show ip flow export
Flow export v9 is enabled for main cache
  Export source and destination details :
  VRF ID : Default
    Source(1)       172.16.1.1 (Unknown)
    Destination(1)  203.0.113.254 (9995)
  Version 9 flow records
  0 flows exported in 0 udp datagrams
  0 flows failed due to lack of export packet
  0 export packets were sent up to process level
  0 export packets were dropped due to no fib
  0 export packets were dropped due to adjacency issues
  0 export packets were dropped due to fragmentation failures
  0 export packets were dropped due to encapsulation fixup
failures
```

Refer to the exhibit. It was noticed that after NetFlow is configured in the router, the collector stopped receiving flow information. Which action resolves the issue?

A. Apply the ip flow egress command to the loopback2 interface.

B. Modify the source through the ip flow-export source loopback1 command.

C. Configure an IP address on the loopback2 interface to use as a source.

D. Change the IP address of the loopback 2 interface to a public IP address.

Correct Answer: *D*

kaupz 1 week ago

Selected Answer: D

tested in the lab, D seems to be the only option, because A & B don't make sense and C disappears the line including "Source" instead of making the source "unknown":

```
R2(config)#do sh ip flow export | inc Source
Source(1) 172.16.1.1 (Loopback2)
R2(config)#no int lo2
R2(config)#do sh ip flow export | inc Source
R2(config)#do sh ip flow export | inc Source
```

upvoted 2 times

kaupz 1 week ago

okay, B is the closest correct answer - but the answer should be "no shut" on lo2

```
R2(config-if)#do sh ip flow export | inc Source
Source(1) 172.16.1.1 (Loopback2)
R2(config-if)#shutdown
R2(config-if)#
000045: *Dec 17 21:12:11: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback2, changed state to down
000046: *Dec 17 21:12:11: %LINK-5-CHANGED: Interface Loopback2, changed state to administratively down
R2(config-if)#
R2(config-if)#do sh ip flow export | inc Source
Source(1) 172.16.1.1 (Unknown)
R2(config-if)#
```

upvoted 2 times

DeWalt95 2 weeks, 2 days ago

Selected Answer: D

D seems to be the only one that makes sense.
The current Lo2 Ip address is within RFC1918 and therefore not routable from/to the public IP destination.

upvoted 1 times

Question #529 *Topic 1*

What is LDP used for in an LSR?

A. to allow for a system-wide exchange of labels across MPLS network

B. to create a label across the PE routers for end-to-end path assignment

C. to communicate the routes known for a specific interface

D. to create a database of label bindings that allow for hop-by-hop forwarding

Correct Answer: *D*

DeWalt95 2 days, 14 hours ago

Selected Answer: D

D is correct.

upvoted 1 times

Question #530 *Topic 1*

What are the two goals of micro BFD sessions? (Choose two.)

A. The high bandwidth member link of a link aggregation group must run BFD.

B. Any member link on a link aggregation group must run BFD.

C. Continuity for each member link of a link aggregation group must be verified.

D. Run the BFD session with 3x3 ms hello timer.

E. Each member link of a link aggregation group must run BFD.

Correct Answer: *CE*

Tedmus 2 weeks, 5 days ago

Selected Answer: CE

Answer is correct:
Reference:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/xe-16-8/irb-xe-16-8-book/irb-micro-bfd.html

The goal of micro BFD sessions are:
- Run BFD session over each LAG member link.
- Verify link continuity for each member link.
- Allow BFD to control the LAG member link to be part of the L2 load-balancing table of the LAG interface in the presence or absence of LACP.

upvoted 2 times

DeWalt95 2 days, 14 hours ago

Thanks for confirming.

upvoted 1 times

Question #531 *Topic 1*

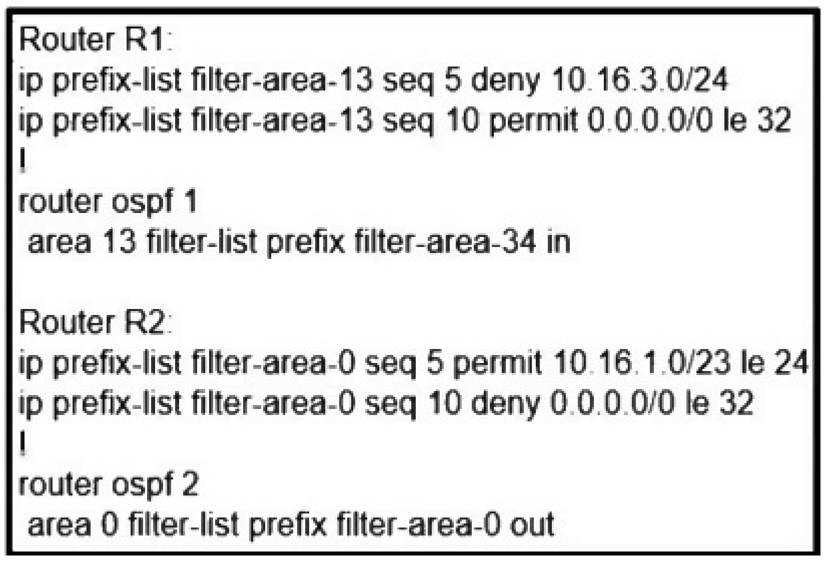

```
Router R1:
ip prefix-list filter-area-13 seq 5 deny 10.16.3.0/24
ip prefix-list filter-area-13 seq 10 permit 0.0.0.0/0 le 32
!
router ospf 1
 area 13 filter-list prefix filter-area-34 in

Router R2:
ip prefix-list filter-area-0 seq 5 permit 10.16.1.0/23 le 24
ip prefix-list filter-area-0 seq 10 deny 0.0.0.0/0 le 32
!
router ospf 2
 area 0 filter-list prefix filter-area-0 out
```

Refer to the exhibit. R1 should receive 10.16.2.0/24 from R2. Which action resolves the issue?

A. Add prefix-list seq 1 on R1 to permit 10.16.0.0/22.

B. Add prefix-list seq 1 on R1 to permit 10.16.2.0/24.

C. Modify prefix-list seq 5 on R2 to permit 10.16.0.0/22.

D. Modify prefix-list seq 5 on R2 to permit 10.16.0.0/23.

Correct Answer: *C*

⊟ **DavideDL** 1 week ago

Selected Answer: C

C is correct , but the prefix-list called in R1 filter list has the wrong name.

upvoted 1 times

⊟ **DeWalt95** 2 weeks, 2 days ago

Selected Answer: C

10.16.1.0/23 is processed as 10.16.0.0/23 (10.16.0.1 -10.16.1.255)

The only prefix that includes the required submit is C

upvoted 1 times

Question #532 *Topic 1*

Which characteristic is representative of a hub-and-spoke topology between PE routers in a Layer 3 MPLS VPN network?

A. The PE routers use different RDs for each VRF to import and export M-BGP prefixes.

B. Each PE router uses a different RD to identify all branches.

C. The PE routers use different RTs to import and export M-BGP prefixes.

D. The PE routers are configured with multiple VRFs for all branches.

Correct Answer: *C*

⊟ **DeWalt95** 2 days, 14 hours ago

Selected Answer: C

C based on carefully parsing the language in the answers

upvoted 1 times

Question #533 Topic 1

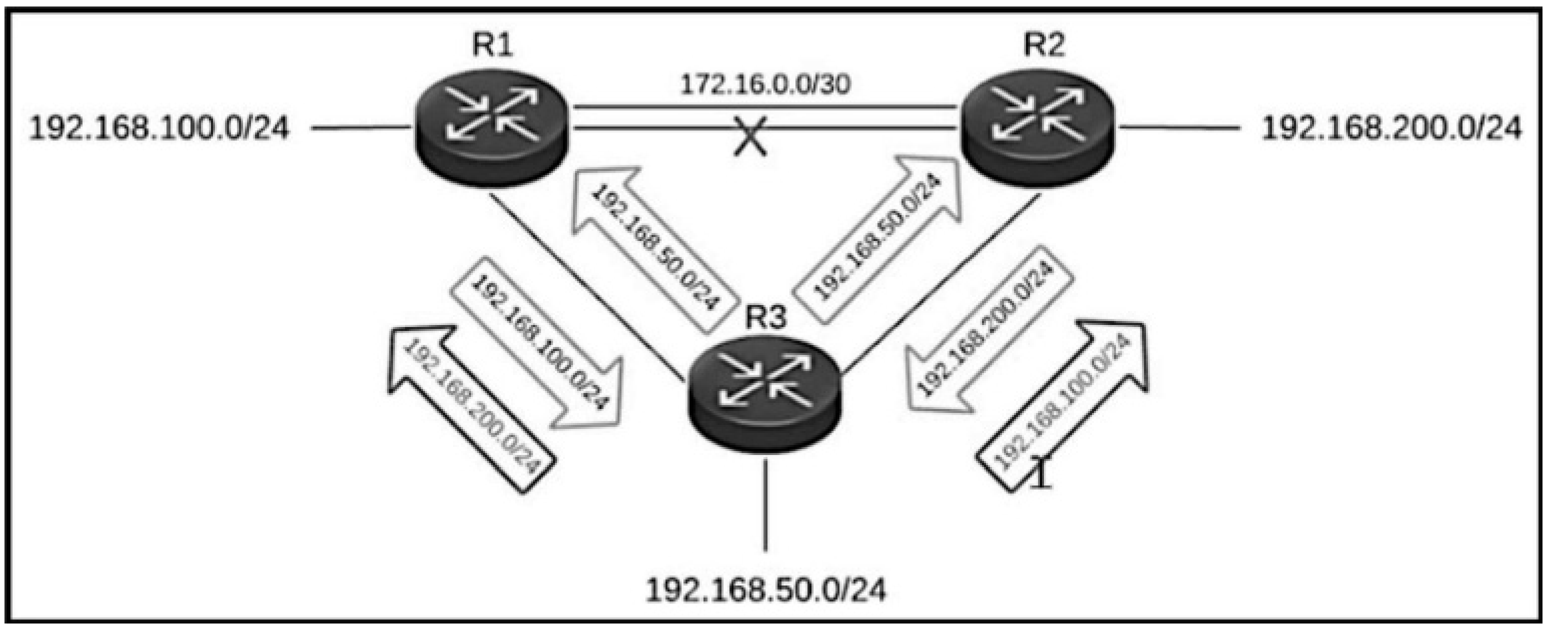

Refer to the exhibit. The primary link between R1 and R2 went down, but R3 is still advertising the 192.168.200.0/24 network to R1 and the 192.168.100.0/24 network to R2, which creates a loop. Which action resolves the issue?

A. Configure the eigrp stub command under the EIGRP process on R2.

B. Configure the summary-address 192.168.0.0 255.255.0.0 100 command on R3.

C. Configure the eigrp stub command under the EIGRP process on R3.

D. Configure the eigrp stub leak-map command under the EIGRP process on R1.

Correct Answer: *C*

⊟ **DeWalt95** 2 weeks, 2 days ago

Not sure where the loop is but to stop R3 being sued for transitive routing between R1+2 making it an EIGRP stub works.

upvoted 1 times

Question #534 Topic 1

How does MPLS Layer 3 VPN function?

A. When an EIGRP internal route is redistributed into BGP by one PE and then back into EIGRP by another PE, the originating router ID for the route is changed to the router ID of the first PE.

B. When a destination PE device receives a labeled packet, it pops the label and uses it to forward the packet to the correct CE device.

C. When a PE device forwards a packet received from a CE device across the provider network, it labels the packet with the label learned from the source PE device.

D. When a VPN route is learned from a CE device and injected into IGP, a VPN route distinguisher attribute is associated with it.

Correct Answer: *B*