Question #1 *Topic 1*

Click the Exhibit button.

```
user@router-re0> show system s?
Possible completions:
     services       Show service applications information
     snapshot       Show snapshot information
     software       Show loaded JUNOS extensions
     statistics     Show statistics for protocol
     storage        Show local storage data
```

You have configured graceful RE switchover (GRES), however you cannot complete the show system switchover command.

Referring to the exhibit, what is the problem?

A. The command is only available if a backup router is configured.

B. The command is only available if non-stop routing is enabled.

C. The command is only available if graceful restart is enabled.

D. The command is only available on the backup Routing Engine.

Correct Answer: *D*

zineeddine Highly Voted 5 months, 2 weeks ago

Issue the show system switchover command only on the backup Routing Engine. This command is not supported on the primary Routing Engine because the kernel-replication process daemon does not run on the primary Routing Engine. This process runs only on the backup Routing Engine.

upvoted 8 times

Question #2 *Topic 1*

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

A. MED

B. origin

C. community

D. AS path

E. next hop

Correct Answer: *BDE*

networkwalker 5 months, 1 week ago

BGP Attribute Categories

There are four categories of BGP attributes:

Well-known mandatory:Recognized by all BGP peers, passed to all peers, and present in all Update messages. Well-known mandatory attributes include:- Next-hop- Origin- AS PATH

Well-known discretionary:Recognized by all routers, passed to all peers, and optionally included in the Update message. Well-known discretionary attributes include:- Local Preference- Atomic Aggregate

Optional transitive:Possibly recognized by BGP routers and passed to BGP peers. Optional transitive attributes are marked as partial when not recognized. Optional transitive attributes include:- Aggregator- Community

Optional non-transitive:Possibly recognized by BGP routers but not passed to peers. Optional non-transitive attributes include:- Multi-exit discriminator (MED)- Originator ID- Cluster-ID

upvoted 3 times

Question #3 *Topic 1*

What is the correct order of BGP attributes for active route selection?

A. next hop –> local preference –> AS path –> MED –> origin

B. next hop –> local preference –> AS path –> origin –> MED

C. next hop –> origin –> local preference –> AS path –> MED

D. next hop –> AS path –> local preference –> origin –> MED

Correct Answer: *B*

RCADENA 8 months ago

B. next hop –> local preference –> AS path –> origin –> MED

upvoted 2 times

Question #4 *Topic 1*

What are two bridging concepts that are used to maintain an Ethernet switching table? (Choose two.)

A. learning

B. exporting

C. aging

D. timing

Correct Answer: *AC*

Sam_2121 10 months ago

Yep, A and C

upvoted 2 times

Zaraoui 10 months, 3 weeks ago

A. learning

C. aging

These concepts are crucial for the efficient operation of an Ethernet switch:

Learning: The switch uses this process to build and update its switching table. When a frame is received, the switch records the source MAC address and the port it arrived on, allowing it to correctly forward future frames to the appropriate port.

Aging: This process ensures that the switching table remains current. Entries that have not been refreshed within a certain period are removed, preventing the table from becoming outdated and freeing up space for new entries

upvoted 2 times

Zaraoui 10 months, 3 weeks ago

A and C

The two bridging concepts used to maintain an Ethernet switching table are:

A. learning: This refers to the process by which a switch gathers information about which devices are connected to which ports. When a frame is received on a port, the switch records the source MAC address and the port it was received on. This information is then stored in the switching table.

C. aging: This refers to the process by which old or stale entries in the switching table are removed. If an entry in the table has not been updated for a certain period, the switch will delete it to ensure that the table remains current and accurate. This helps in maintaining efficient and up-to-date network communication.

upvoted 1 times

Question #5 *Topic 1*

Which BGP attribute is used to detect routing loops?

A. AS path

B. MED

C. local preference

D. next hop

Correct Answer: *A*

Zaraoui 10 months, 3 weeks ago

A. AS path

AS path is a mandatory attribute in BGP that contains the sequence of AS numbers that the route has traversed. It helps prevent routing loops by ensuring that a BGP speaker does not accept a route that contains its own AS number in the AS path attribute. This prevents a BGP speaker from advertising a route back to the AS from which it learned the route, thereby avoiding routing loops in the BGP network.

upvoted 2 times

Question #6 *Topic 1*

How does a Junos device learn about MAC addresses when it is first connected to an Ethernet LAN?

A. The device sends out a network multicast message asking for all devices and MAC addresses on the network and stores this information in addition to the interface from which the response was received.

B. The device sends out a network broadcast message asking for all devices and MAC addresses on the network and stores this information in addition to the interface from which the response was received.

C. The device learns the destination MAC addresses from traffic in the network and stores this MAC address in addition to the interface from which the traffic was received.

D. The device learns the source MAC addresses from traffic in the network and stores this MAC address in addition to the interface from which the traffic was received.

Correct Answer: *D*

Zaraoui 10 months, 3 weeks ago

D. The device learns the source MAC addresses from traffic in the network and stores this MAC address in addition to the interface from which the traffic was received.

Here's how this process typically works:

When a Junos device is connected to an Ethernet LAN, it listens to the traffic on each of its interfaces.
As traffic flows through the device, it observes the source MAC addresses of the frames arriving on each interface.
The Junos device then associates each MAC address it observes with the specific interface from which it received the traffic.
This information is stored in the Ethernet switching table (also known as MAC address table) of the device.
By observing the source MAC addresses in the incoming frames, the device builds its understanding of which MAC addresses are reachable via which interfaces.
upvoted 2 times

Question #7 *Topic 1*

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

A. By default, an IS-IS interface has a maximum metric value of 1023.

B. Wide metrics enable IS-IS interface to have a maximum metric value of around 16 million.

C. Wide metrics enable an IS-IS interface to have a maximum metric value of 1023.

D. By default, an IS-IS interface has a maximum metric value of 63.

Correct Answer: *BD*

Community vote distribution

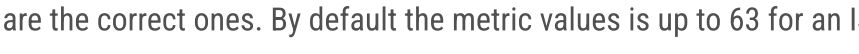

desertprick 7 months, 2 weeks ago

Selected Answer: BD

BD - The default TOTAL PATH metric for IS-IS is 1023, but the maximum metric for an interface is 63

upvoted 2 times

desertprick 7 months, 2 weeks ago

BD - The default TOTAL PATH metric for IS-IS is 1023, but the maximum metric for an interface is 63

upvoted 1 times

42e0062 9 months ago

Selected Answer: BD

By default, the maximum metric value for an IS-IS interface is 63. This is because the original prefix and topology TLVs (IP Internal Reachability, IP External Reachability, and IS reachability) only gave 6 bits to the metric value.

With wide metrics, the cost of an interface can range from 1 to 16,777,215.

upvoted 2 times

dimits 12 months ago

Selected Answer: BD

BD are the correct ones. By default the metric values is up to 63 for an ISIS enabled interface (this is why 1023 is the default max metric for any ISIS path.....the metric can't go much higher is we calculate based on interfaces using non-wide values (up to 63), so adding multiple less-than-63 interfaces can't result easy in a summary more than 1023.

upvoted 1 times

eoj8 1 year, 5 months ago

Selected Answer: D

BD should be the answer

upvoted 1 times

c946f3e 1 year, 7 months ago

Selected Answer: BD

BD CORRECT

upvoted 3 times

Sbiari 1 year, 7 months ago

Selected Answer: BD

Maximum metric value 63

wide metrics enable an is-is interface to have a maximum metric value of around 16 million

upvoted 3 times

SH_Lee 1 year, 8 months ago

Selected Answer: BD

I think correct, B,D

upvoted 4 times

zineeddine 1 year, 11 months ago

it should be BD :
metric—Metric value.

Range: 1 through 63, or 1 through 16,777,215 (if you have configured wide metrics)

Default: 10 (for all interfaces except lo0), 0 (for the lo0 interface)

upvoted 4 times

dennis1988 1 year, 10 months ago

Correct.
Normally, IS-IS metrics can have values up to 63. The total cost to a destination is the sum of the metrics on all outgoing interfaces along a particular path from the source to the destination. By default, the total path metric is limited to 1023.

upvoted 1 times

Sailor 1 year, 11 months ago

All IS-IS routes have a cost, which is a routing metric that is used in the IS-IS link-state calculation. The cost is an arbitrary, dimensionless integer that can be from 1 through 63, or from 1 through 16,777,215 (224 – 1) if you are using wide metrics.

upvoted 3 times

Question #8 *Topic 1*

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

A. IBGP neighbors must use the same AS number.

B. By default, the TTL on product-related packets for external neighbors is set to 1.

C. EBGP neighbors must use the same AS number.

D. By default, the TTL on protocol-related packets for internal neighbors is set to 1.

Correct Answer: *AB*

Community vote distribution

StefBrakki 1 month, 1 week ago

Selected Answer: AB

AB. Explanation:

Option A:

IBGP (Internal BGP) neighbors operate within the same Autonomous System (AS). Therefore, all IBGP neighbors must share the same AS number to establish a proper IBGP session.

Option B:

The TTL (Time-to-Live) value for BGP packets sent to external BGP (EBGP) neighbors is set to 1 by default. This ensures that EBGP sessions are typically established between directly connected routers. For longer-distance EBGP connections, you can modify the TTL using the "ebgp-multihop" command.

upvoted 1 times

Epyon 1 month, 2 weeks ago

Selected Answer: AB

Per Juniper's own documentation, default TTL for eBGP is 1 for directly-connected single hop and 64 for multi-hop. The default TTL for iBGP is 255. These are the same values that all router vendors use.

upvoted 1 times

2e6bc5f 4 months, 2 weeks ago

Selected Answer: AD

TTL for iBGP neighbours I set to 1 cos iBGP peers are typically directly connected or expected to be within the same network

upvoted 1 times

ztw3587t 7 months, 2 weeks ago

Selected Answer: AB

A and B

eBGP TTL is1

iBGP TTL is 64

upvoted 1 times

desertprick 7 months, 2 weeks ago

Selected Answer: AD

default Time-to-Live (TTL) value for BGP TCP connections to internal neighbors is 1, which prevents packets from being forwarded beyond the local AS.

upvoted 1 times

Zaraoui 10 months, 1 week ago

Selected Answer: AB

A and B

upvoted 2 times

tnh24 1 year ago

By default, the TTL (Time-to-Live) on protocol-related packets (such as BGP TCP connections) for internal neighbors (within the same AS) is set to 1. This default TTL value helps prevent BGP packets from being inadvertently forwarded beyond the local AS.

upvoted 1 times

Question #9 *Topic 1*

Click the Exhibit button.

```
[edit]
user@switch# show interfaces ge-0/0/1
native-vlan-id 20;
unit 0 {
     family ethernet-switching {
          interface-mode trunk;
          vlan {
               members [ 10 20 ];
          }
     }
}
```

Referring to the exhibit, what will happen to untagged frames?

A. The untagged frames are associated with VLAN 10.

B. The untagged frames are dropped.

C. The untagged frames are load balanced between VLAN 10 and VLAN 20.

D. The untagged frames are associated with VLAN 20.

Correct Answer: *D*

Community vote distribution

D (100%)

Zaraoui 10 months, 3 weeks ago

Selected Answer: D

D. The untagged frames are associated with VLAN 20.

upvoted 2 times

Question #10 *Topic 1*

An OSPF router does not have a router ID configured.
In this scenario, which statement is correct about the router ID?

A. The Junos OS will use the IP address assigned to the interface with the lowest MAC address.

B. The Junos OS will use the IP address assigned to the loopback interface for the router ID.

C. A router ID will not be assigned until it is manually configured.

D. The Junos OS will use the IP address assigned to the interface with the highest priority.

Correct Answer: *B*

Community vote distribution

B (100%)

zineeddine Highly Voted 1 year, 11 months ago

The router identifier is used by BGP and OSPF to identify the routing device from which a packet originated. The router identifier usually is the IP address of the local routing device. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used.

upvoted 8 times

Zaraoui Most Recent 10 months, 3 weeks ago

Selected Answer: B

the B is correct

upvoted 1 times

Question #11 *Topic 1*

You are asked to configure filter-based forwarding on a Junos device.
Which two statements are correct in this scenario? (Choose two.)

A. You must create a routing policy.

B. You must create a route target.

C. You must create and apply a match filter.

D. You must create a routing instance.

Correct Answer: *CD*

Community vote distribution

Sailor Highly Voted 1 year, 11 months ago

Selected Answer: CD

To configure filter-based forwarding, perform the following tasks:
Create a match filter on the ingress device. To specify a match filter, include the filter filter-name statement at the [edit firewall] hierarchy level. A packet that passes through the filter is compared against a set of rules to classify it and to determine its membership in a set. Once classified, the packet is forwarded to a routing table specified in the accept action in the filter description language. The routing table then forwards the packet to the next hop that corresponds to the destination address entry in the table.
Create routing instances that specify the routing table(s) to which a packet is forwarded, and the destination to which the packet is forwarded at the [edit routing-instances] hierarchy level.

upvoted 5 times

maj Most Recent 2 weeks, 1 day ago

Selected Answer: AC

The two correct statements when configuring filter-based forwarding (FBF) on a Junos device are:

A. You must create a routing policy.

C. You must create and apply a match filter.

Here's why:

A. You must create a routing policy: Filter-based forwarding relies on routing policies to define the criteria (the filter) for selecting a specific forwarding next hop or routing instance. The policy dictates what traffic will be treated specially.

C. You must create and apply a match filter: The core of FBF is identifying specific traffic flows based on certain characteristics (source/destination IP address, port numbers, protocol, etc.). This identification is done using a firewall filter (which acts as the "match filter"). This filter defines the criteria that, when matched, will trigger the forwarding action defined in the routing policy. The filter must then be applied to an interface to inspect traffic.

upvoted 1 times

Takigama 7 months, 1 week ago

Selected Answer: CD

However, technically you don't need an RI to do filter based forwarding as you can make the filter send to either an interface or an IP as its next hop...

upvoted 1 times

dimits 12 months ago

Selected Answer: CD

Filter policy and routing instance should be the correct!

upvoted 1 times

h2michael 1 year, 5 months ago

Selected Answer: CD

Filter-based forwarding (FBF), which is also called Policy Based Routing (PBR), provides a a simple but powerful way to route IP traffic to different interfaces on the basis of Layer-3 or Layer-4 parameters.

FBF works by using match conditions in a firewall filter to select certain traffic and then direct it to a given routing instance that points to the desired next hop. To ensure the next hop is resolvable, interface routes from the main routing table are shared via RIB group with the routing table(s) specified in the routing instance(s).

Match conditions can include the source or destination IP address, source or destination port, IP protocol, DSCP value, TCP flag, ICMP type, and packet length.

upvoted 1 times

Sbiari 1 year, 7 months ago

Selected Answer: CD

You must creat and apply a match filter

you must creat a routing instance

upvoted 2 times

Laminejuve 1 year, 11 months ago

Selected Answer: CD

FBF works by using match conditions in a firewall filter to select certain traffic and then direct it to a given routing instance that points to the desired next hop. To ensure the next hop is resolvable, interface routes from the main routing table are shared via RIB group with the routing table(s) specified in the routing instance(s).

upvoted 2 times

Laminejuve 1 year, 11 months ago

FBF works by using match conditions in a firewall filter to select certain traffic and then direct it to a given routing instance that points to the desired next hop. To ensure the next hop is resolvable, interface routes from the main routing table are shared via RIB group with the routing table(s) specified in the routing instance(s).

upvoted 1 times

Question #12 *Topic 1*

You are implementing traffic engineering in your MPLS network. You must ensure that the MPLS routes are used to traverse your network. Your solution should not affect IGP routes in your route tables.
In this scenario, which traffic engineering setting will accomplish this behavior?

A. bgp-igp-both-ribs

B. mpls-forwarding

C. bgp-igp

D. bgp

Correct Answer: *D*

Community vote distribution

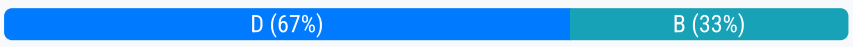

Sailor Highly Voted 1 year, 11 months ago

Selected Answer: D

bgp—On BGP destinations only. Ingress routes are installed in the inet.3 routing table.
bgp-igp—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 routing table. If IGP shortcuts are enabled, the shortcut routes are automatically installed in the inet.0 routing table.
bgp-igp-both-ribs—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 and inet.3 routing tables. This option is used to support VPNs.
mpls-forwarding—On both BGP and IGP destinations. Use ingress routes for forwarding only, not for routing.
upvoted 8 times

Darkhorsejsboo Highly Voted 1 year, 3 months ago

The correct answer is B. mpls-forwarding. This setting ensures that the MPLS routes are used to traverse your network without affecting IGP routes in your route tables. It enables MPLS traffic engineering tunnels on an interface1. This is a key part of configuring MPLS Traffic Engineering (TE), which allows network operators to define how traffic flows through an MPLS core network2. The other options (A, C, and D) are related to BGP and IGP routing, not specifically to MPLS traffic engineering. Please note that actual configuration may vary depending on the specific network devices and software versions you are using. Always refer to the official documentation for your hardware and software for the most accurate information.
upvoted 6 times

dhikra 12 months ago
i agree with you answer B
upvoted 1 times

StefBrakki Most Recent 1 month, 1 week ago

Selected Answer: B

The correct answer is: B. mpls-forwarding
Explanation:
The mpls-forwarding setting ensures that MPLS routes are used to forward traffic across your MPLS network. It enables traffic to traverse MPLS Label-Switched Paths (LSPs) for traffic engineering purposes.
This setting does not interfere with or alter the IGP routes in your routing tables. IGP routes are still used for control plane operations, while MPLS-forwarding directs traffic through the LSPs in the data plane.

Why not other options?
A. bgp-igp-both-ribs: This setting merges BGP and IGP routes into both the BGP and IGP routing tables, which can affect IGP routes, contrary to the requirements.
C. bgp-igp: This setting installs BGP routes into the IGP routing table, which again would impact IGP routes.
D. bgp: This applies only to BGP behavior and does not specifically address MPLS traffic engineering.
upvoted 1 times

Epyon 1 month, 2 weeks ago

Selected Answer: B

Per Juniper docs: "To use LSPs for forwarding but exclude them from route selection, include the mpls-forwarding option for the traffic-engineering statement." This would leave BGP and iBGP unaffected.

upvoted 1 times

2e6bc5f 4 months, 2 weeks ago

Selected Answer: B

MPLS forwarding is the answer, it ensures it uses LSPs where applicable

upvoted 1 times

yhmj 5 months ago

Selected Answer: B

The mpls-forwarding setting allows MPLS to take precedence for forwarding traffic without altering the IGP routes in the routing table.

upvoted 1 times

zineeddine 1 year, 11 months ago

sould be answer D

upvoted 3 times

Question #13 *Topic 1*

Click the Exhibit button.

```
user@switch> show spanning-tree bridge
STP bridge parameters
Context ID                              : 0
Enabled protocol                        : RSTP
     Root ID                            : 8192.59.c5:8d:ae:db:41
     Hello time                         : 10 seconds
     Maximum age                        : 40 seconds
     Forward delay                      : 30 seconds
     Message age                        : 0
     Number of topology changes         : 6
     Time since last topology change    : 781 seconds
     Topology change initiator          : ge-0/0/14.0
     Topology change last recvd. from   : 2c:6b:f5:31:06:0b
     Local parameters
          Bridge ID                     : 8192.50:c5:8d:ae:db:41
          Extended system ID            : 0
          Internal instance ID          : 0
```

Which two statements are correct about the information shown in the exhibit? (Choose two.)

A. The switch has a bridge priority of 8k.

B. This switch is the root bridge for this spanning tree topology.

C. The root bridge is reachable using the ge-0/0/14 interface.

D. The root bridge's priority is 4k.

Correct Answer: *AC*

Community vote distribution

AC (74%) AB (26%)

StefBrakki 1 month, 1 week ago

Selected Answer: AC

A. This switch has a bridge priority of 8k.

C. The root bridge is reachable using the ge-0/0/14 interface.

Explanation:

Option A: The exhibit indicates that the switch's bridge priority is 8k (8192). This is part of the switch's Bridge ID.

Option C: The root bridge is indeed reachable via the ge-0/0/14 interface, as shown in the spanning-tree information.

Why Option B is Incorrect:

Since the Root ID and Bridge ID differ, this switch is not the root bridge.

upvoted 1 times

Epyon 1 month, 2 weeks ago

Selected Answer: AC

In this exhibit, the bridge IDs are different. Moral of the story based on comments: pay very close attention to the bridge ID's when taking the test.

upvoted 1 times

2e6bc5f 4 months, 2 weeks ago

Selected Answer: AC

Different bridge IDs so it's not the root bridge

upvoted 1 times

DinoVercotti 5 months, 3 weeks ago

Selected Answer: AB

There is a typo in the exhibit, the Root ID and the Bridge ID is same.

Both are 8192.50:c5:8d:ae:db:41

upvoted 2 times

ztw3587t 7 months, 2 weeks ago

Selected Answer: AC

A and C are correct.

A - Bridge priority is 8k (8192) -> true
B - Root ID and Bridge ID are the same -> false
C - Root bridge is reachable by ge-0/0/14 interface -> true
D - Bridge priority is 4k (4096) -> false
upvoted 1 times

neighborhoodgorilla 1 year ago

Selected Answer: AC

AC is the answer, root id is different from bridge id, also the priority is not 4k, 8192 is 8k
upvoted 1 times

Sailor 1 year, 11 months ago

Selected Answer: AC

Root ID is different from Bridge ID
upvoted 2 times

Sailor 1 year, 11 months ago

Selected Answer: AB

Switch@juniper>show spanning-tree bridge

STP bridge parameters
Context ID : 0
Enabled protocol : RSTP
Root ID : 32768.28:c0:da:3d:50:40 <----------
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Message age : 0
Number of topology changes : 2
Time since last topology change : 3218 seconds
Topology change initiator : ge-0/0/4.0
Topology change last recvd. from : 2c:6b:f5:8b:23:04
Local parameters
Bridge ID : 32768.28:c0:da:3d:50:40 <----------
Extended system ID : 0
Internal instance ID : 0

If values of the Root ID and the Bridge ID have equal values as in the above example, then we confirm that this is the Root Bridge or the Root Switch.

Root ID Bridge ID of the elected spanning-tree root bridge. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.
upvoted 3 times

antigel8 1 year, 2 months ago

The root ID and bridge ID aren't the same in the exhibit.
upvoted 3 times

networkwalker 1 year, 11 months ago

Selected Answer: AC

The bridge ID and root ID are not the same.
upvoted 3 times

zineeddine 1 year, 11 months ago

should be AC
upvoted 3 times

Laminejuve 1 year, 11 months ago

Selected Answer: AC

this bridge is not the root bridge
upvoted 4 times

Question #14 *Topic 1*

Which statement is true about the BGP active state?

A. The BGP active state is the initial state where all BGP traffic is refused.

B. The BGP active state is when BGP attempts to acquire a peer by initiating a TCP connection.

C. The BGP active state is when BGP waits for the TCP connection to be established.

D. The BGP active state is when BGP exchanges update, notification, and keepalive messages with its peer.

Correct Answer: *B*

Community vote distribution

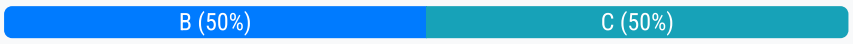

StefBrakki 1 month, 1 week ago

Selected Answer: C

In the Active state, BGP has already initiated a TCP connection with the peer (during the Connect state), but it is now waiting for the TCP connection to be fully established. If the connection cannot be established, BGP retries or transitions back to the Connect state.

upvoted 1 times

Zaraoui 10 months, 3 weeks ago

Selected Answer: B

C is incorrect. Waiting for the TCP connection to be established is part of the Connect state in BGP, not the Active state. The Connect state follows the Active state once a TCP connection attempt is made.

upvoted 1 times

Question #15 *Topic 1*

Click the Exhibit button.

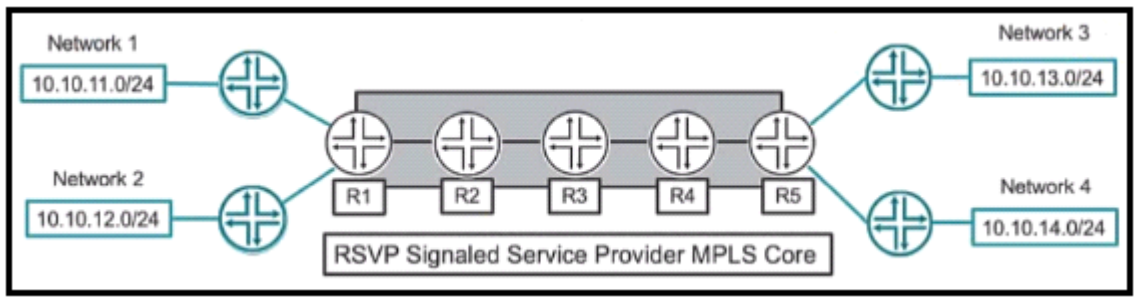

Referring to the exhibit, what is the minimum number of LSPs required to support all four networks?

A. 1

B. 2

C. 8

D. 4

Correct Answer: *B*

Community vote distribution

B (80%) A (20%)

loadme Highly Voted 1 year, 11 months ago

The correct answer should B

upvoted 8 times

loadme Highly Voted 1 year, 11 months ago

It is B, you only need "2"

upvoted 7 times

ztw3587t Most Recent 7 months, 2 weeks ago

Selected Answer: B

LSP R1 to R5

LSP R5 to R1

upvoted 1 times

42e0062 8 months, 3 weeks ago

Selected Answer: A

RSVP Signaled Service Provider MPLS Core is shown in the example. Only a single Label Switched Path (LSP) is required to support all four networks because MPLS can use label stacking to multiplex multiple LSPs over a single physical path. One LSP can be established through the MPLS core, and all networks can use this LSP with different label stacks

upvoted 1 times

Zaraoui 10 months, 1 week ago

Selected Answer: B

only two : lsb-R1-to-R5 and lsb-R5-to-R1

upvoted 1 times

Saran_Br 1 year, 1 month ago

Selected Answer: B

R1 to R8 are MPLS CORE Routers,

LSP1 for R1 to R8,

LSP2 for R8 to R1

upvoted 2 times

samchakr 1 year, 9 months ago

It is C. RSVP LSP from -

Network 1 to Network 3

Network 1 to Network 4

Network 2 to Network 3

Network 2 to Network 4
and vice-versa
upvoted 3 times

Question #16 *Topic 1*

Which two protocols are capable of distributing labels for segment routing? (Choose two.)

A. RSVP

B. LDP

C. IS-IS

D. OSPF

Correct Answer: *CD*

Community vote distribution

CD (83%) AB (17%)

ADharma 10 months, 1 week ago

Selected Answer: CD

I used IS-IS and OSPF for my segment routing at work..

upvoted 3 times

Sailor 11 months, 1 week ago

Selected Answer: CD

Laminejuve is right, my bad

upvoted 2 times

loadme 11 months, 1 week ago

The question here is about segment routing it's not about what is used to build MPLS.

We all know that in normal MPLS network, you will have a combination of LDP to create the label bindings, you will also have IGP as underlay (OSPF, IS-IS) to route your internal traffic and BGP as overlay that will describe your customer information. Basically you have three protocols to build your MPLS network IGP, LDP, BGP.

Now segment routing allows us to simply this by getting rid of the unnecessary protocols like LDP, RSVP etc. and use our IGPs (OSPF, IS-IS) to distribute labels bindings. This is to reduce loads on the control plane.

I'm 100% sure the CORRECT ANSWER IS: CD

upvoted 3 times

Sailor 11 months, 2 weeks ago

Selected Answer: AB

LDP and RSVP is used to build MPLS LSPs

upvoted 1 times

Laminejuve 11 months, 2 weeks ago

segment routing isn't a protocol in the same way that RSVP or LDP are protocols. There is no "segment routing protocol" that two routers talk to each other to exchange MPLS labels.

Instead, the labels are advertised directly inside IS-IS or OSPF!

upvoted 4 times

Question #17 *Topic 1*

Which two steps are required to enable MPLS on a physical interface in Junos? (Choose two.)

A. Add family mpls on the interface.

B. Add the loopback interface under protocols mpls.

C. Add family mpls on the loopback interface.

D. Add the interface under protocols mpls.

Correct Answer: *AD*

Community vote distribution

AD (100%)

Zaraoui 10 months, 3 weeks ago

Selected Answer: AD

To enable MPLS on a physical interface in Junos, the correct steps are:

A. Add family mpls on the interface.
D. Add the interface under protocols mpls.

Here's an explanation of each step:

A. Add family mpls on the interface: This command configures the physical interface to support MPLS. It enables MPLS forwarding on the interface and allows MPLS labels to be processed for packets received on that interface.

D. Add the interface under protocols mpls: This step associates the physical interface with the MPLS protocol. It allows the interface to participate in MPLS signaling and forwarding by configuring it under the MPLS protocol configuration.

upvoted 1 times

c946f3e 1 year, 6 months ago

1
2
3
4
5
6
set interfaces ge-0/0/1.0 family mpls
set interfaces ge-0/0/2.0 family mpls
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface ge-0/0/2.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0

upvoted 1 times

Click the Exhibit button.

```
[edit]
user@R1# show interfaces ae0
aggregated-ether-options {
     minimum-links 1;
     lacp {
          passive;
          periodic fast;
     }
}
unit 0 {
     family inet {
          address 10.0.0.1/30;
     }
}
user@R1> show interfaces terse | match ae0
ge-0/0/1.0     up   up   aenet          --> ae0.0
ge-0/0/2.0     up   up   aenet          --> ae0.0
ae0            up   down
ae0.0          up   down inet           10.0.0.1/30
user@R1> show lacp statistics interfaces
Aggregated interface: ae0
     LACP Statistics:    LACP Rx   LACP Tx   Unknown Rx     Illegal Rx
          ge-0/0/1            0         0              0              0
          ge-0/0/2            0         0              0              0
user@R2# show interfaces ae0
aggregated-ether-options {
     minimum-links 2;
     lacp {
          passive;
          periodic slow;
     }
}
unit 0 {
     family inet {
          address 10.0.0.2/30;
     }
}
```

You configured an aggregated Ethernet bundle between R1 and R2. Unfortunately, the bundle is not working correctly.

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

A. The LACP minimum links configuration should be changed to 2 on R1.

B. The LACP periodic mode should be changed to fast on R2.

C. The LACP mode should be changed to active on R1.

D. The LACP minimum links configuration should be changed to 1 on R2.

Correct Answer: *C*

Community vote distribution

C (75%) A (25%)

wassupkay 6 months ago

C - passive means both are waiting for a connect to establish LACP session. thus, neither are sending the connect, so none will establish LAC.

upvoted 1 times

Rbrahmi 7 months, 1 week ago

Selected Answer: C

The correct answer is C

upvoted 1 times

GolanZemach 7 months, 3 weeks ago

Selected Answer: C

Answer C is correct at least one speaker needs to be in active mode

upvoted 2 times

GolanZemach 7 months, 4 weeks ago

Answer C is correct at least one speaker needs to be in active mode

upvoted 1 times

antigel8 8 months, 3 weeks ago

Selected Answer: A

In R1, ge-0/0/1 and 2 are attached to ae0 which is configured with a minimum-link of 1. R2 ae configure has a minimum of 2 which is correct. R1 needs to increase the minimum of links to 2.

upvoted 1 times

antigel8 8 months, 3 weeks ago

I am wrong in my previous statement. Answer is to turn at least one speaker to active mode

upvoted 2 times

Question #19 *Topic 1*

What are two requirements for a unified in-service software upgrade? (Choose two.)

A. Bidirectional Forwarding Detection must be enabled on the device.

B. Nonstop active routing must be enabled on the device.

C. The device must have dual Routing Engines.

D. The device must be part of the chassis cluster.

Correct Answer: *BC*

Community vote distribution

BC (100%)

Zaraoui 10 months, 3 weeks ago

Selected Answer: BC

The unified in-service software upgrade (ISSU) feature enables you to upgrade your device between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU is supported only on dual Routing Engine platforms. In addition, the graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) features must be enabled.

upvoted 1 times

Question #20 *Topic 1*

By default, which two statements are correct about switch ports on a Junos device? (Choose two.)

A. Trunk ports receive and transmit untagged traffic.

B. Access ports receive and transmit tagged traffic.

C. Trunk ports receive and transmit tagged traffic.

D. Access ports receive and transmit untagged traffic.

Correct Answer: *CD*

Community vote distribution

CD (100%)

Zaraoui 10 months, 3 weeks ago

Selected Answer: CD

C and D

upvoted 1 times

eoj8 1 year, 6 months ago

That is correct!

upvoted 1 times

c946f3e 1 year, 7 months ago

Selected Answer: CD

vlan's tagging on the Trunk Port

vlan untagging on the Access Port

upvoted 3 times

Question #21 *Topic 1*

You want to share routes between two routing instances that you have configured?
What are two ways to accomplish this task? (Choose two.)

A. Use a non-forwarding instance.

B. Configure an instance import policy

C. Create a forwarding instance.

D. Use a RIB group.

Correct Answer: *BD*

Community vote distribution

Sailor Highly Voted 1 year, 5 months ago

Selected Answer: BD

the methods for communicating between routing instances are as follows:

static route with a next-hop of next-table pointing to the appropriate routing table which contains more accurate information
rib-groups to mirror routing information from one route-table to another. However, in many cases, in order to make this work, interface-routes also need to be mirrored. RIB Group policy can be used to constrain the routing information
instance-import and instance-export statements configured within the individual routing-instances to leak routes from one table to another. Again, policy can be used here to constrain the routing information. This method is more straightforward than the rib-group method
A final approach is to use physical interfaces or logical-tunnels to stitch routing-instances and use a routing protocol or static routes across this connection between the two routing-instances.

upvoted 7 times

Laminejuve 1 year, 5 months ago

agree so BD will be the correct answers

upvoted 4 times

alicenaihi Most Recent 10 months ago

Selected Answer: BD

BD will be the correct

upvoted 1 times

Bob_Smith 1 year, 1 month ago

should be BD...Create the RIB group first, then filter on what is actually imported with an import policy

upvoted 2 times

Misho_e_pich_2023 1 year, 1 month ago

Selected Answer: BD

A. Use a non-forwarding instance: A non-forwarding instance is used for tasks that require routing information but do not participate in the actual packet forwarding. It does not directly facilitate the sharing of routes between routing instances.

C. Create a forwarding instance: A forwarding instance is used for packet forwarding in Junos OS. While it's involved in the forwarding of packets, it's not a mechanism for sharing routes between routing instances.

So, options B and D are the correct ways to share routes between routing instances in Junos OS.

upvoted 2 times

zineeddine 1 year, 5 months ago

should be BD

upvoted 3 times

Question #22 *Topic 1*

Click the Exhibit button.

```
user@R1> show configuration protocols bgp
group IBGP_CORE {
     type internal;
     local-address 192.168.1.1;
     export NEXT_HOP_SELF;
     neighbor 192.168.1.5;
}
group CUSTOMER_A {
     type external;
     peer-as 65101;
     neighbor 10.1.101.101;
}
user@R1> show configuration policy-options
policy-statements NEXT_HOP_SELF {
     term SELF {
          then {
               next-hop self;
          }
     }
}
```

Which statement is correct about the configuration shown in the exhibit?

A. This policy will reject all prefixes because it does not end with a then accept statement.

B. The NEXT_HOP_SELF policy must also be applied to the group called CUSTOMER_A for bidirectional communication to succeed.

C. This router will take prefixes learned from external BGP and change the BGP next hop to itself before re-advertising them to IBGP peers.

D. This router will take prefixes learned from internal BGP and change the BGP next hop to itself before re-advertising them to EBGP peers.

Correct Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #23 *Topic 1*

Click the Exhibit button.

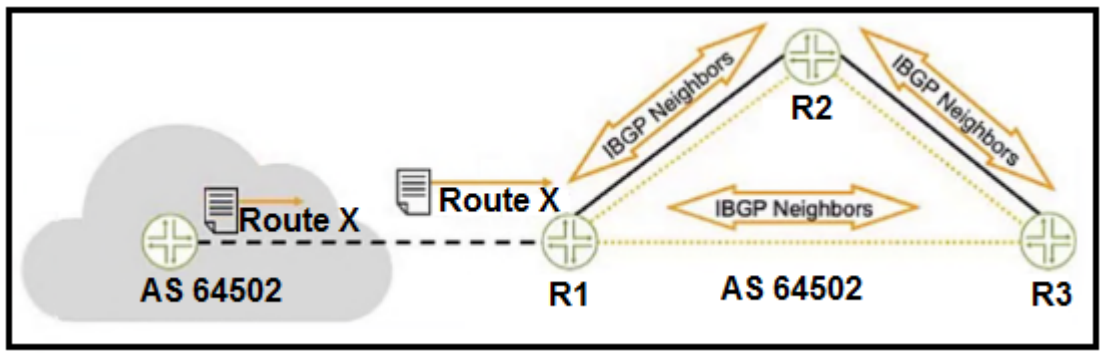

Referring to the exhibit, from which device(s) does R3 learn about Route X?

A. both R2 and R1

B. R1 only

C. directly from the router in AS 64502

D. R2 only

Correct Answer: *B*

Community vote distribution

B (100%)

42e0062 9 months ago

This diagram doesn't make sense. Route X appears to be coming from a router in the same AS as R1,2,and 3. In which case unless R1 is configured as a router reflector none of the other routers beyond R1 would learn of route x.

upvoted 1 times

ADharma 11 months, 4 weeks ago

Selected Answer: B

R1 only

upvoted 1 times

alicenaihi 1 year, 4 months ago

Selected Answer: B

R1 only

upvoted 2 times

h2michael 1 year, 5 months ago

Selected Answer: B

R2 cannot advertise IBGP learned routes to another IBGP peers

upvoted 1 times

Bob_Smith 1 year, 7 months ago

B

R2 will not forward IBGP learned routes to another IBGP peer

upvoted 2 times

Bhupi1979 1 year, 9 months ago

Selected Answer: B

B is correct

upvoted 3 times

Bhupi1979 1 year, 9 months ago

B is correct

upvoted 3 times

Farman09 1 year, 9 months ago

Selected Answer: B

B is correct

upvoted 2 times

ay_dos 1 year, 11 months ago

expect it is a typo, all Routers are in same AS. no valid answer

upvoted 2 times

Luxury179 1 year, 11 months ago

Selected Answer: B

B should be the correct answer

upvoted 2 times

zineeddine 1 year, 11 months ago

from R1 only due to loop prevention.

upvoted 2 times

Laminejuve 1 year, 11 months ago

Selected Answer: B

R2 can not forward IBGP learned routes to another IBGP neighbor correct

upvoted 2 times

Sailor 1 year, 11 months ago

Selected Answer: B

R2 can not forward IBGP learned routes to another IBGP neighbor to prevent loops

upvoted 2 times

Question #24 *Topic 1*

Click the Exhibit button.

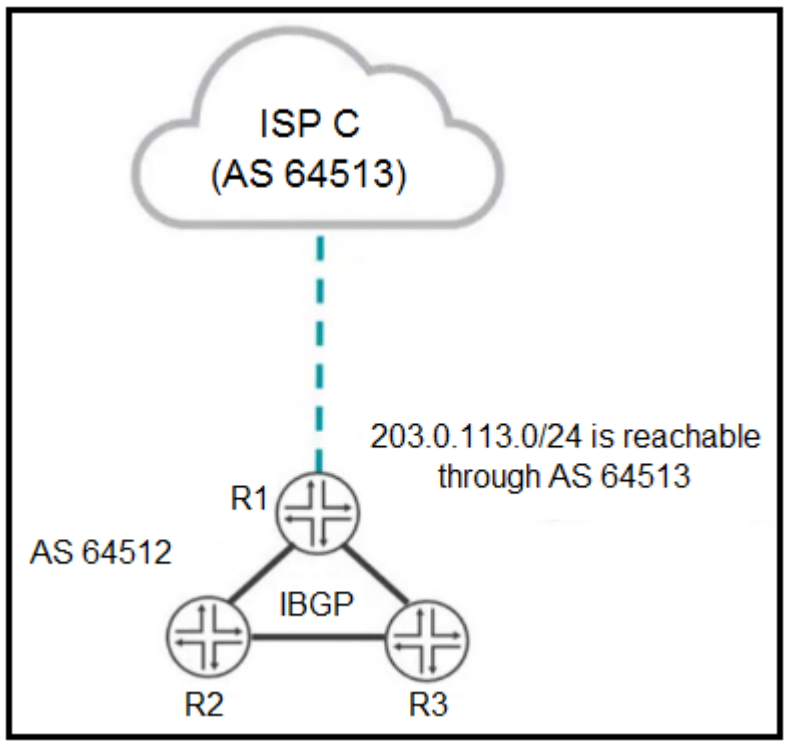

You confirm that the R2 and R3 routers are receiving a BGP route to the 203 0.113.0/24 network, but both routers display the route as hidden.

Referring to the exhibit, which two actions solve this problem? (Choose two.)

A. Apply the routing policy on R1 as an import policy to the IBGP group.

B. Configure a routing policy on R1 that sets the next hop for the 203 0.113.0/24 BGP route to the IP address that R1 uses for IBGP peering.

C. Configure a routing policy on R1 that sets the next hop for the 203.0.113.0/24 BGP route to the IP address that R1 uses for EBGP peering.

D. Apply the routing policy on R1 as an export policy to the IBGP group.

Correct Answer: *BD*

Community vote distribution

BD (100%)

alicenaihi 10 months ago

Selected Answer: BD

Selected Answer: BD

upvoted 2 times

oz1i 11 months, 3 weeks ago

CD, carefull

upvoted 1 times

teresoj352 1 year, 2 months ago

Selected Answer: BD

B. Configure a routing policy on R1 that sets the next hop for the 203 0.113.0/24 BGP route to the IP address that R1 uses for IBGP peering. = configure a "then next-hop self" policy

D. Apply the routing policy on R1 as an export policy to the IBGP group. = export that to ur iBGP peers

so, B, D are correct, so, in other words, whoever picked B and C either isn't a network engineer, either is some certifications freak who just wants to have as many certs as possible and doesn't give a damn about the real world, and the trade of being a real network engineer, why I say that? cuz makes no darn sense, if you do as 'suggested' you'll be making some policies no-one uses as they're not applied anywhere.....

upvoted 3 times

networkwalker 1 year, 5 months ago

Selected Answer: BD

in the export policy, set as next hop self

upvoted 2 times

zineeddine 1 year, 5 months ago

i agree BD

upvoted 1 times

Sailor 1 year, 5 months ago

Selected Answer: BD

apply the policy to IBGP group to send it to IBGP neighbors

upvoted 2 times

Laminejuve 1 year, 5 months ago

yes i agree you have to configure a policy to change the next hop and then apply it as an exposrt policy to the IBGP peers so BD are the correct answers

upvoted 1 times

Question #25 *Topic 1*

Click the Exhibit button.

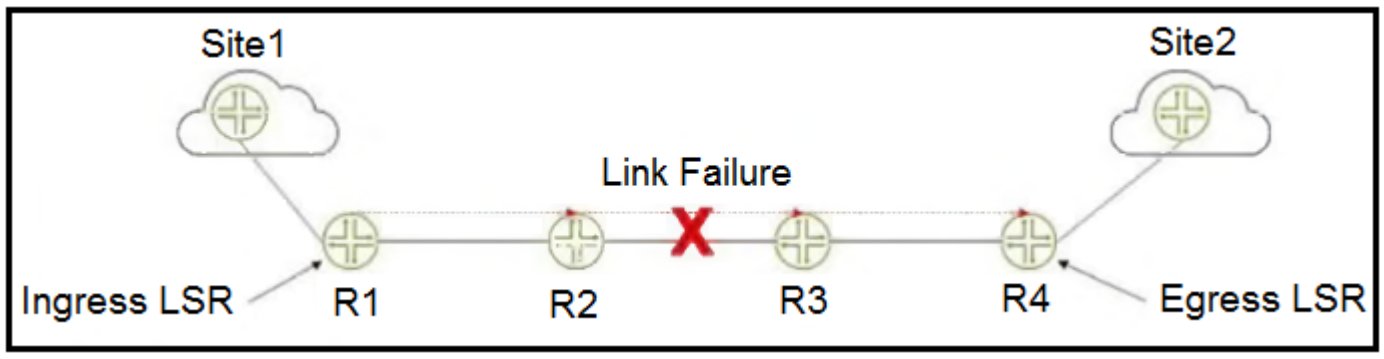

Referring to the exhibit, you have an established RSVP LSP between R1 and R4 when you experience a link failure between R2 and R3.
Which two statements are correct? (Choose two.)

A. R2 sends a ResvTear message upstream to R1 signaling the link failure.

B. R3 sends a PathTear message downstream to R4 signaling the link failure.

C. R2 sends a PathTear message upstream to R1 signaling the link failure.

D. R3 sends a ResvTear message downstream to R4 signaling the link failure.

Correct Answer: *AB*

Community vote distribution

CD (100%)

StefBrakki 1 month, 1 week ago

Selected Answer: CD

C. R2 sends a PathTear message upstream to R1 signaling the link failure.

D. R3 sends a ResvTear message downstream to R4 signaling the link failure.

Explanation:

PathTear (Option C):

When the link between R2 and R3 fails, R2 detects the failure and sends a PathTear message upstream to R1. This informs the upstream nodes about the broken path so they can take corrective actions.

ResvTear (Option D):

Similarly, R3 sends a ResvTear message downstream to R4. This notifies the downstream nodes that the reserved resources for the LSP on the failed link are no longer valid, and they should release them.

upvoted 1 times

nmcgruder 1 month, 1 week ago

https://www.juniper.net/documentation/us/en/software/junos/mpls/topics/topic-map/rsvp-overview.html#id-rsvp-message-types__d145436e556

RSVP Overview from Juniper shows RevTear goes back upstream to the sending router.

upvoted 2 times

zineeddine 5 months, 2 weeks ago

agree AB

upvoted 2 times

Question #26 *Topic 1*

Click the Exhibit button.

```
[edit]
user@switch# show interfaces
ge-1/0/0 {
     vlan-tagging;
     encapsulation flexible-ethernet-services;
     unit 100 {
          encapsulation vlan-bridge;
          vlan-id 100;
          input-vlan-map {
               push;
               vlan-id 200;
          }
          output-vlan-map pop;
     }
}
ge-1/0/4 {
     stacked-vlan-tagging;
     encapsulation flexible-ethernet-services;
     unit 0 {
          encapsulation vlan-bridge;
          vlan-tags outer 200 inner 100;
     }
}
…
[edit]
user@switch# show bridge-domains
customer1 {
     interface ge-1/0/0.100;
     interface ge-1/0/4.0;
}
```

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

A. Traffic ingressing ge-1/0/0 that is tagged with VLAN 200 will be dropped.

B. Traffic ingressing ge-1/0/0 that is tagged with VLAN 100 will be dropped.

C. Traffic ingressing ge-1/0/0 that is tagged with VLAN 100 will egress ge-1/0/4 with an outer VLAN tag of 200.

D. Traffic ingressing ge-1/0/0 that is tagged with VLAN 200 will egress ge-1/0/4 with an outer VLAN tag of 200.

Correct Answer: *AC*

Currently there are no comments in this discussion, be the first to comment!

Question #27 *Topic 1*

Which two interface types are used as tunnel endpoints? (Choose two.)

A. ge

B. ae

C. gr

D. ip

Correct Answer: *CD*

Community vote distribution

CD (100%)

desertprick 5 months ago

Selected Answer: CD

In Juniper Networks, the two interface types used as tunnel endpoints are the internal interfaces gr-0/0/0 and ip-0/0/0. These interfaces are used for Generic Routing Encapsulation (GRE) and IP-IP tunnels on SRX Series Firewalls. The Junos OS creates these interfaces at system bootup, and they are not associated with physical interfaces

upvoted 1 times

c946f3e 1 year, 6 months ago

tunnel-end-point name {
ipv4 {
source-address 10.255.1.1;
destination-address 10.255.2.0/25;
}
gre {
key 9;
}
}

https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/ref/statement/tunnel-end-point-edit-firewall.html

upvoted 1 times

roychoar 1 year, 6 months ago

"endpoint", IMHO, means an interface that can be the source and/or the destination of a tunnel interface. I would have chosen A and B

upvoted 1 times

Question #28 *Topic 1*

Click the Exhibit button.

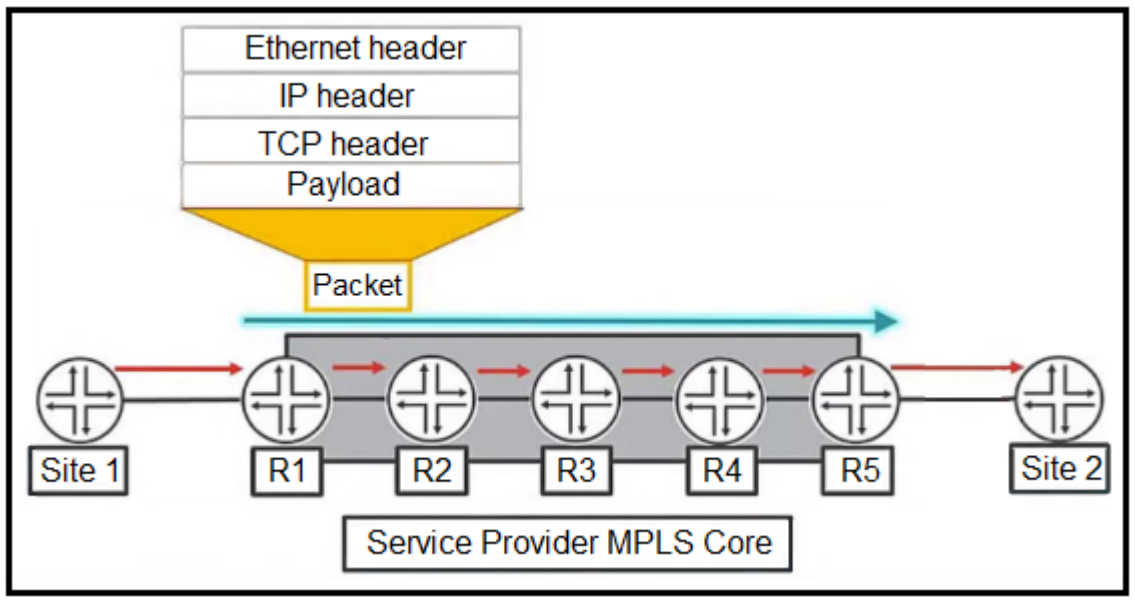

Which two statements are correct about the actions taken as the packet traverses the service provider MPLS network from Site1 to Site2 as shown in the exhibit? (Choose two.)

A. R1 will perform a lookup using the inet.3 table.

B. R1 will perform a lookup using the mpls.0 table.

C. R2 will perform a lookup using the mpls.0 table.

D. R2 will perform a lookup using the inet.3 table.

Correct Answer: *AC*

Currently there are no comments in this discussion, be the first to comment!

Question #29 *Topic 1*

Click the Exhibit button.

```
[edit]
user-@R1# show protocols mpls
label-switched-path R1-to-R6 {
     to 10.1.1.6;
     secondary via-R2-R4;
     secondary any-path;
}
path via-R2-R4 {
     10.1.1.2 strict;
     10.1.1.4 strict;
}
path any-path;
interface ge-0/0/0.0;
interface ge-0/0/1.0;
```

All devices in the network are configured correctly and the path requirements are valid.

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

A. The secondary LSP using the via-R2-R4 path will not be signaled, and its state will be down.

B. The secondary LSP using the any-path path will be signaled, and its state will be up.

C. The secondary LSP using the any-path path will not be signaled, and its state will be down.

D. The secondary LSP using the via-R2-R4 path will be signaled, and its state will be up.

Correct Answer: *CD*

Community vote distribution

CD (60%) AC (40%)

Ak_baba 7 months ago

Selected Answer: AC

When the primary LSP is up, secondary static LSPs are not signaled and are in a "down" state. To instruct a router to pre-signal a secondary LSP, you can configure it with the keyword "standby". Without this keyword, the router will only signal the secondary path after the primary path fails.

https://saidvandeklundert.net/2015-05-13-juniper-primary-and-secondary-lsp/

upvoted 1 times

Ak_baba 7 months ago

Sorry, I missed something.

In this case, the primary LSP is not configured. If you do not configure a primary path, the first secondary path that is established is selected as the path.

Answer: CD

https://www.juniper.net/documentation/us/en/software/junos/mpls/topics/topic-map/primary-secondary-static-lsp-configuration.html

upvoted 1 times

42e0062 8 months, 2 weeks ago

Selected Answer: CD

When you have 2 secondary paths, they will be signaled in the order listed in the config. So LSP R1-to-R6, only the "via-R2-R4" will be computed. The "any-path will" be in a down state.

upvoted 1 times

Zaraoui 10 months, 1 week ago

Selected Answer: AC

the secondary lsp path will be never signaled. will be signaled only when the keyword "standby " is added.

upvoted 3 times

antigel8 1 year, 2 months ago

Selected Answer: CD

I think C and D are the valids

upvoted 2 times

dydzah 1 year, 4 months ago

Selected Answer: CD

only one path is signaled at a time,

upvoted 3 times

Question #30 *Topic 1*

Click the Exhibit button.

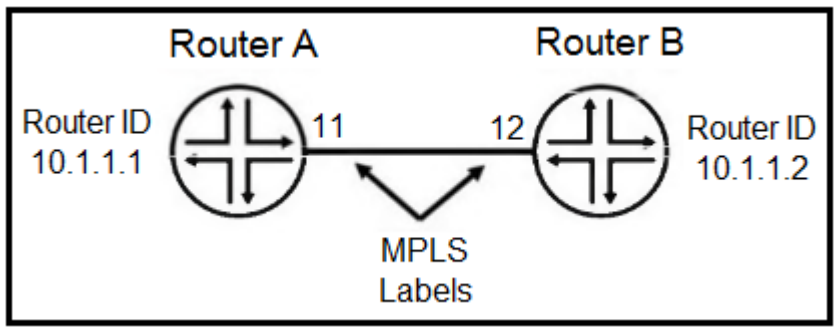

The routers shown in the exhibit are configured for segment routing.

In this scenario, what is the adjacency SID that Router B advertises to Router A?

A. 11

B. 10.1.1.2

C. 10.1.1.1

D. 12

Correct Answer: *D*

Community vote distribution

D (73%) B (18%) 9%

zineeddine Highly Voted 1 year, 11 months ago

Shloud be answer B

Currently, Junos OS enables you to configure a SPRING node SID for IPv4 and IPv6 address families for each routing instance. This node SID is attached to an IPv4 and IPv6 router ID if the router ID is configured on the loopback interface. Otherwise, the lowest IP address assigned to the loopback interface is chosen as the node SID. Configuring a node SID through policy allows you to choose the loopback address that gets the node SID. If the node SID configuration exists and a policy is defined for node SID selection for the same prefix, then the policy configuration takes precedence.

upvoted 7 times

42e0062 Most Recent 8 months, 2 weeks ago

Selected Answer: D

As others have indicated. The adjacency SID that Router B advertises to Router A in a segment is determined in this scenario by the MPLS label indicated in the exhibit. so the Answer is D. I don't know why there are so many upvotes on the wrong answer.....

upvoted 2 times

julin_10 1 year, 2 months ago

Selected Answer: D

The adjacency SID in segment routing is advertised by one router to inform another about the label that should be used to forward traffic to it over a particular link. In this scenario, Router B would advertise an adjacency SID to Router A, which would be usedby Router A to forward traffic to Router B. Based on the exhibit, the adjacency SID that Router B advertises to Router A would be 12.

References: Juniper Networks documentation on Segment Routing: Segment Routing Overview

upvoted 3 times

Jayjit 1 year, 3 months ago

Correct answer is D

https://packetpushers.net/blog/yet-another-blog-about-segment-routing-part-1/

upvoted 1 times

Darkhorsejsboo 1 year, 3 months ago

The adjacency SID (Segment Identifier) that Router B advertises to Router A in a segment routing configuration is indicated by the MPLS label associated with the link as seen from Router B's perspective. According to the diagram, Router B is associated with the MPLS label 12 for the link between Router A and Router B.

Therefore, the adjacency SID that Router B advertises to Router A is:

D. 12

upvoted 3 times

theeprosemaestro 1 year, 3 months ago

Answer is D. 12. Adjacency SIDs are used in segment routing for MPLS using labels. The answer is 12 (D).

upvoted 3 times

alicenaihi 1 year, 4 months ago

Selected Answer: D

Adjacency SID is an MPLS label that is generated and advertised in IS-IS or OSPF for each adjacency.

It is 12 in this question

upvoted 3 times

7c3a129 1 year, 4 months ago

Selected Answer: A

should b letter A because image mentioned that labels are 11 and 12 because adjancency SID are dinamically assigned. It is not related with loopback(node SID). So finally if some device wants send packets to Router B must has label 11 in stack. This label was advertised by Router B in IGP.

upvoted 1 times

[Removed] 1 year, 4 months ago

Selected Answer: B

zineeddine Highly Voted 6 months, 2 weeks ago

Shloud be answer B

Currently, Junos OS enables you to configure a SPRING node SID for IPv4 and IPv6 address families for each routing instance. This node SID is attached to an IPv4 and IPv6 router ID if the router ID is configured on the loopback interface. Otherwise, the lowest IP address assigned to the loopback interface is chosen as the node SID. Configuring a node SID through policy allows you to choose the loopback address that gets the node SID. If the node SID configuration exists and a policy is defined for node SID selection for the same prefix, then the policy configuration takes precedence.

upvoted 2 times

Question #31 *Topic 1*

Which OSPF database packet determines which router is in charge of the database synchronization and the transferring of LSA headers between the two systems?

A. link-state request

B. database description

C. hello

D. link-state update

Correct Answer: *B*

Community vote distribution

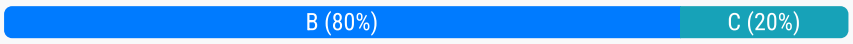

zineeddine Highly Voted 1 year, 11 months ago

should be B

upvoted 8 times

StefBrakki Most Recent 1 month, 1 week ago

Selected Answer: B

B. database description

Explanation:

The Database Description (DD) packet in OSPF is used during the database synchronization process between two OSPF routers. It provides an overview of the router's Link State Advertisements (LSAs) by including their headers.

During the exchange, one router acts as the master, and the other as the slave. The master is in charge of controlling the synchronization process.

The DD packets help in determining which LSAs need to be requested or updated, ensuring both routers have consistent link-state databases.

upvoted 1 times

42e0062 8 months, 2 weeks ago

Selected Answer: B

For the link state routing protocols, it is essential that the link state database for all the routers remain synchronized. This synchronization will begin as soon as an adjacency formed between the neighbors. OSPF uses DBD - database descriptor packets for that purpose.

upvoted 1 times

julin_10 1 year, 2 months ago

Selected Answer: B

The Database Description (DD) packets serve two main purposes:

1. determining which router is in charge of the database synchronization

2. transferring the LSA headers between the two systems

The Database Description (DBD) OSPF packet is used during the adjacency formation process to describe the contents of the topological database between routers. The routers exchange DBD packets to determine the master and slave relationship, which dictates the router in charge of the database synchronization.

References: OSPF Design Guide, Juniper Networks Documentation

OSPF Database Description Packets, Juniper Networks Documentation

upvoted 2 times

[Removed] 1 year, 4 months ago

Selected Answer: B

Correct answer is B

upvoted 1 times

rankapan17 1 year, 8 months ago

Selected Answer: B

the Database Description (DD) packets serve two main purposes:

1. determinining which router is in charge of the database synchronization

2. tansferring the LSA headers between the two systems

upvoted 2 times

RC_123 1 year, 9 months ago

Selected Answer: B

B is correct

upvoted 3 times

RC_123 1 year, 8 months ago

Changing answer to C

upvoted 1 times

RC_123 1 year, 9 months ago

B is correct

upvoted 1 times

RC_123 1 year, 8 months ago

Changing answer to C

upvoted 1 times

blondeman 1 year, 9 months ago

Selected Answer: B

https://ipwithease.com/ospf-packet-types/

At the time of adjacency is being initialized, these packets are exchanged. These packets describe topological database contents. The database may be described by using multiple packets. A poll-response procedure is used for the description of multiple packets usage. Among the routers, one is designated to be master, and the other a slave. The Database Description packets are sent by the slave after sending the Database Description packets by the master.

upvoted 2 times

blondeman 1 year, 9 months ago

The answer is B.

https://ipwithease.com/ospf-packet-types/

upvoted 1 times

mrnipsnips 1 year, 10 months ago

Selected Answer: C

Hello packet

When several systems or routers that run OSPF have interfaces attached to a common network, the Hello protocol determines the designated router (DR) .

upvoted 3 times

Question #32 *Topic 1*

Click the Exhibit button.

```
user@R2> show ospf route
Topology default Route Table:
Prefix              Path      Route     NH    Metric    NextHop        NextHop
                    Type      Type      Type            Interface      addr/label
192.168.1.1         Intra     AS BR     IP         1    ge-0/0/3.0     172.26.1.1
192.168.1.3         Intra     Area BR   IP         1    ge-0/0/1.0     172.26.2.2
172.18.1.0/24       Ext2      Network   IP         0    ge-0/0/3.0     172.26.1.1
172.26.1.0/30       Intra     Network   IP         1    ge-0/0/3.0
172.26.2.0/30       Intra     Network   IP         1    ge-0/0/1.0
172.26.3.0/30       Intra     Network   IP       100    ge-0/0/2.0
172.26.4.0/30       Inter     Network   IP         2    ge-0/0/1.0     172.26.2.2
192.168.1.1/32      Ext2      Network   IP         1    ge-0/0/3.0     172.26.1.1
192.168.1.2/32      Intra     Network   IP         0    lo0.0
192.168.1.3/32      Intra     Network   IP         1    ge-0/0/1.0     172.26.2.2
192.168.1.4/32      Inter     Network   IP         2    ge-0/0/1.0     172.26.2.2
```

Which prefix in the output shown in the exhibit is an external prefix injected by an OSPF router?

A. 192.168.1.3

B. 172.18.1.0/24

C. 192.168.1.4

D. 172.26.4.0/30

Correct Answer: *B*

Currently there are no comments in this discussion, be the first to comment!