Question #1 *Topic 1*

You are asked to automatically provision new Juniper Networks devices in your network with minimal manual intervention.
Before you begin, which two statements are correct? (Choose two.)

A. You must have a system log (syslog) server to manage system log messages and alerts.

B. You must have a DHCP server that provides the location of the software image and configuration files.

C. You must have a file server that stores software image and configuration files.

D. You must have an NTP server to perform time synchronization.

**Suggested Answer:** *BC*

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

Question #2 *Topic 1*

You are asked for TX and RX traffic statistics for each interface to which an application server is attached. The statistics need to be reported every five seconds.
Using the Junos default settings, which telemetry method would accomplish this request?

A. Native Sensors

B. OpenConfig

C. SNMP

D. gNMI

**Suggested Answer:** *A*

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

Question #3 *Topic 1*

You are asked to deploy 100 QFX Series devices using ZTP. Each QFX5120 requires a different configuration.
In this scenario, what are two components that you would configure on the DHCP server? (Choose two.)

A. the IP address of the FTP server

B. the MAC address for each QFX5120

C. the management IP address for each QFX5120

D. the MAC address of the FTP server.

**Suggested Answer:** *AC*

*Community vote distribution*

AB (100%)

**5298a69** 2 weeks, 3 days ago

**Selected Answer: AB**

A.its necesary B.)the mac address of QFX each device to assign a specific IP address

upvoted 1 times

Question #4 *Topic 2*

Which two statements are true about IP fabrics using unnumbered BGP? (Choose two.)

A. Unnumbered BGP peering automatically provisions IPv4 peering.

B. Unnumbered BGP requires that family inet6 is configured on each interface.

C. Unnumbered BGP peering automatically provisions IPv6 peering.

D. Unnumbered BGP requires that family inet is configured on each interface.

**Suggested Answer:** *CD*

*Community vote distribution*

BC (100%)

**svregaz** 2 months ago

**Selected Answer: BC**

It's to provision IPv6 and it does need an inet6 ip address

upvoted 1 times

**MichelT** 4 months ago

**Selected Answer: BC**

See https://www.juniper.net/documentation/us/en/software/nce/nce-225-bgp-unnumbered/index.html

upvoted 1 times

Click the Exhibit button.

```
{master:0}[edit]
user@leaf1# show policy-options
...
policy-statement load-balance {
    term 1 {
        then {
            load-balance per-packet;
        }
    }
}
{master:0}[edit]
user@leaf1# show routing-options
router-id 192.168.100.11;
autonomous-system 65100;
{master:0}[edit]
user@leaf1# show protocols
bgp {
    group spine {
        type external;
        export direct;
        local-as 65003;
        multipath {
            multiple-as;
        }
        neighbor 172.16.1.5 {
            peer-as 65001;
        }
        neighbor 172.16.1.17 {
            peer-as 65002;
        }
    }
}
```

You are troubleshooting an IP fabric for your data center. You notice that your traffic is not being load balancing to your spine devices from leaf devices.

Referring to the configuration shown in the exhibit, what must be configured to solve this issue?

A. The load-balance policy must have a from statement that matches on protocol bgp.

B. The multipath multiple-as configuration must be configured for each peer in the BGP spine group.

C. The load-balance policy must be applied as an export policy to your BGP spine group.

D. The load-balance policy must be applied to the forwarding table under the routing-options hierarchy.

**Suggested Answer:** *D*

*Community vote distribution*


**svregaz** 2 months ago

**Selected Answer: D**

routing-options {
forwarding-table {
export LOAD-BALANCE-POLICY;
}
}

upvoted 1 times

**MichelT** 4 months ago

**Selected Answer: D**

normal ECMP policy

upvoted 1 times

**niciov71** 4 months, 2 weeks ago

**Selected Answer: D**

The load-balance policy must be applied as export policy of forwarding-table under routing-options hierarchy for ECMP traffic load balancing

upvoted 1 times

Question #6 *Topic 2*

You are deploying a Clos IP fabric with an oversubscription ratio of 3:1.
In this scenario, which two statements are correct? (Choose two.)

A. The oversubscription ratio decreases when you add spine devices.

B. The oversubscription ratio increases when you remove spine devices.

C. The oversubscription ratio remains the same when you add spine devices.

D. The oversubscription ratio remains the same when you remove spine devices.

**Suggested Answer:** *BC*

*Community vote distribution*

AB (100%)

**svregaz** 2 months ago

**Selected Answer: AB**

Spines command the oversub ratio, so it is logical that adding some it will decrease the oversubscription whilst remove some it will increase it

upvoted 1 times

**MichelT** 4 months ago

**Selected Answer: AB**

changing the number of spines allways has impact on the oversubscription.

upvoted 1 times

Question #7 *Topic 2*

You are asked to set up an IP fabric that supports AI or ML workloads. You have chosen to use lossless Ethernet.
In this scenario, which statement is correct about congestion management?

A. The switch experiencing the congestion notifies the source device.

B. Only the source and destination devices need ECN enabled.

C. ECN marks packets based on WRED settings.

D. ECN is negotiated only among the switches that make up the IP fabric for each queue.

**Suggested Answer:** *A*

*Community vote distribution*

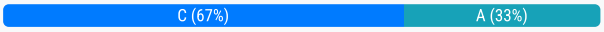


**Moscow** 4 weeks ago

**Selected Answer: C**

not A: ...receiver echoes the congestion notification to the sender

upvoted 1 times

**Kornrono** 3 months ago

**Selected Answer: C**

From the same reference as "MichelT" suggestion, but the correct should be "C" based on this wording

"ECN-enabled devices determine the queue congestion state based on the WRED packet drop profile configuration applied to the queue, so each ECN-enabled queue must also have a WRED drop profile. "

and the component that notify source is "Receiver", based on this wording
" Receiver echoes the congestion notification to the sender."

upvoted 1 times

**MichelT** 4 months ago

**Selected Answer: A**

See https://www.juniper.net/documentation/us/en/software/junos/traffic-mgmt-qfx/cos/topics/concept/cos-qfx-series-explicit-congestion-notification-understanding.html

upvoted 1 times

Question #8 *Topic 2*

You are designing an IP fabric for a large data center, and you are concerned about growth scalability.
Which two actions would you take to address these concerns? (Choose two.)

A. Use EX4300 Series devices as the spine devices.

B. Design a three-stage Clos IP fabric.

C. USE QFX5700 Series devices as the super spines.

D. Design a five-stage Clos IP fabric.

**Suggested Answer:** *BC*

*Community vote distribution*

CD (100%)

**svregaz** 2 months ago

**Selected Answer: CD**

Using a 5 stage model in larger DC also allow to split the DC in smaller pods or portions, simplifying the management of the DC

upvoted 1 times

**MichelT** 4 months ago

**Selected Answer: CD**

5 stage can handle lager traffic volumes then 3 stage

upvoted 1 times

Question #9 *Topic 2*

Which two statements are correct about an IP fabric? (Choose two.)

A. EBGP is only required to route host routing information to external devices outside the fabric.

B. All leaf devices can use the same AS number in an IP fabric without making any adjustments to the EBGP configuration.

C. The multipath multiple-as statement is required to enable ECMP if every device has a different AS number.

D. Only a single point to point EBGP session is required between peers in an IP fabric.

**Suggested Answer:** *AC*

*Community vote distribution*

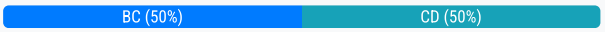


**basvdw107** 2 months, 2 weeks ago

**Selected Answer: CD**

two correct statements are:
C. The multipath multiple-as statement is required to enable ECMP if every device has a different AS number.
D. Only a single point to point EBGP session is required between peers in an IP fabric.
Statement C is correct because in an IP fabric where each device uses a unique AS number, the "multipath multiple-as" statement is essential to enable ECMP (Equal-Cost Multi-Path). Without this configuration, BGP would typically only consider paths from the same AS for load balancing.
Statement D is correct because IP fabrics use a simple connectivity model where only one EBGP session is needed between any two devices (like between a leaf and spine). Additional sessions would be unnecessary and inefficient.
The other statements are incorrect:

Statement A is wrong because EBGP is used for internal fabric connectivity, not just external routing
Statement B is wrong because using the same AS number on all leaf devices would require special configurations to prevent route rejection
upvoted 1 times

**Kornrono** 3 months ago

**Selected Answer: BC**

It is possible for all leaf devices in an IP fabric to use the same AS number. In fact, this is a common configuration, particularly in designs where all leaf devices within a particular POD or area share the same routing policy and AS.

In a Same-Tier-AS IP fabric, all leaf devices can typically share the same AS number without requiring adjustments to the EBGP configuration. In this design, all spine switches use a different, unique AS number, and all leaf switches also share a unique AS number distinct from the spine switches.

https://www.juniper.net/documentation/us/en/software/nce/sg-005-data-center-fabric/topics/task/5-stage-ip-fabric-cloud-dc-configuring.html

https://www.cisco.com/c/en/us/td/docs/dcn/ndfc/1221/articles/ndfc-managing-bgp-based-routed-fabrics/managing-bgp-based-routed-fabrics.html
upvoted 1 times

Question #10 *Topic 2*

You are deploying an IP fabric using EBGP and notice that your leaf devices are advertising and receiving all the routes. However, the routes are not installed in the routing table and are marked as hidden.
Which two statements describe how to solve the issue? (Choose two.)

A. You need to configure a next-hop self policy.

B. You need to configure multipath multiple-as.

C. You need to configure loops 2.

D. You need to configure as-override.

**Suggested Answer:** *AD*

*Community vote distribution*

BC (100%)

**Kornrono** 3 months ago

**Selected Answer: BC**

B. Because,
In an IP fabric deployment using eBGP, configuring multipath with the multiple AS option is crucial to enable Equal Cost Multi-Path (ECMP) load balancing. This ensures traffic is distributed across multiple paths with equal cost, enhancing redundancy and performance. If all devices in the fabric use the same AS number, you don't need the "multiple AS" option.

C. Because,
when deploying an IP fabric using eBGP, you will need to configure loopback addresses on each device. This is because eBGP relies on direct connectivity, and loopback addresses provide a stable and reachable IP address for peering between devices. Loopbacks ensure that BGP sessions can be established and maintained, even if the physical links between devices.
upvoted 1 times

Question #11 *Topic 2*

You are adding a server to a tenant's network within your data center and must limit access to a specific traffic type within the tenant network without pushing all tenant traffic through a firewall.
What will satisfy this requirement?

A. Use filter-based forwarding.

B. Put the new server on a unique subnet within the tenant's network.

C. Use route leaking with EVPN and a routing policy.

D. Use a static route in the tenant VRF with a firewall as the next hop for traffic to the new server.

**Suggested Answer:** *A*

*Community vote distribution*

D (100%)

**5298a69** 4 days, 13 hours ago

**Selected Answer: D**

A. its imposible because its complex to manage at scale and not tenant-specific

upvoted 1 times

Question #12 *Topic 2*

You are deploying a new network to support your AI workloads on devices that support at least 400 Gbps Ethernet. There is no requirement for any Layer 2 VLANs in this network.
Which network architecture would satisfy this requirement?

A. an IP fabric using the EVPN-MPLS architecture

B. an IP fabric with an EVPN-VXLAN architecture

C. an IP fabric using the PIM-SM to signal VXLAN overlay

D. an IP fabric using EBGP

**Suggested Answer:** *B*

*Community vote distribution*

D (100%)

**basvdw107** 3 months, 1 week ago

**Selected Answer: D**

The correct answer is D. an IP fabric using EBGP.
Let me explain why:
For AI workloads requiring high-speed Ethernet (400 Gbps) without any Layer 2 VLAN requirements, you need a scalable, high-performance Layer 3 network architecture. An IP fabric using EBGP is the most appropriate choice because:

It's a pure Layer 3 solution that doesn't waste resources on unnecessary Layer 2 overlay technologies
EBGP provides excellent scalability and fast convergence
It's less complex than the other options since you don't need any overlay protocols

The other options are less suitable because:

Options A and B (EVPN-MPLS and EVPN-VXLAN) are primarily designed to extend Layer 2 connectivity across a Layer 3 fabric, which is unnecessary given your requirements
Option C (PIM-SM with VXLAN) is also focused on providing Layer 2 services, which aren't needed

Since your AI workloads only require high-throughput Layer 3 connectivity without VLANs, the simpler IP fabric with EBGP provides the most efficient and appropriate architecture..
upvoted 3 times

Question #13 *Topic 2*

You want to convert an MX Series router from a VXLAN Layer 2 gateway to a VXLAN Layer 3 gateway for VNI 100. You have already configured an IRB interface.
In this scenario, which command would you use to accomplish this task?

A. set bridge-domains VLAN-100 routing-interface irb.100

B. set protocols isis interface irb.100 passive

C. set protocols ospf area 0.0.0.0 interface irb.100 passive

D. set vlans VLAN-100 13-interface irb.100

**Suggested Answer:** *A*

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

Question #14 *Topic 2*

You are asked to identify microburst traffic occurring in the network leading to packet drops in your data center switches.
Which two tools would be used in this scenario? (Choose two.)

A. port buffer monitoring

B. port mirroring

C. syslog

D. Traceoptions

**Suggested Answer:** *AB*

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

Question #15 *Topic 2*

You are deploying an IP fabric with an oversubscription ratio of 3:1.
In this scenario, which two statements are correct? (Choose two.)

A. The oversubscription ratio remains the same when you remove leaf devices.

B. The oversubscription ratio increases when you remove leaf devices.

C. The oversubscription ratio decreases when you add leaf devices.

D. The oversubscription ratio remains the same when you add leaf devices.

**Suggested Answer:** *BC*

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

Question #16 *Topic 2*

You are deploying multiple Juniper switches at the same location. Your switches are currently using the factory-default configuration.
In this scenario, which two statements are correct? (Choose two.)

A. the DHCP server configuration cannot provide Junos version requirements to DHCP clients.

B. The switch will try to request an IP address from a DHCP server using all interfaces that are connected and are operational.

C. The switch will try to request an IP address from a DHCP server using only the management interface.

D. The DHCP server configuration can provide Junos version requirements to DHCP clients.

**Suggested Answer:** *BC*

*Community vote distribution*

BD (100%)

**basvdw107** 3 months, 1 week ago

**Selected Answer: BD**

The two correct statements about Juniper switches with factory-default configuration are:
B. The switch will try to request an IP address from a DHCP server using all interfaces that are connected and are operational.
D. The DHCP server configuration can provide Junos version requirements to DHCP clients.
Let me explain why these are correct:
B. When a Juniper switch is in factory-default configuration, it will attempt to obtain an IP address via DHCP across all operational interfaces, not just the management port. This is part of Juniper's Zero Touch Provisioning (ZTP) functionality, allowing the switch to get connectivity through any available path. This increases the chances of successful auto-provisioning.
D. Juniper's DHCP implementation supports vendor-specific options that allow the DHCP server to communicate Junos version requirements to the switches. This is an important part of the Zero Touch Provisioning process, as it allows centralized control of OS versions across network devices.
upvoted 2 times

Question #17 *Topic 2*

Which two statements are true about a pure IP fabric? (Choose two.)

A. An IP fabric does not support Layer 2 protocols.

B. An IP fabric supports Layer 2 VLANs.

C. Devices in an IP fabric function as Layer 3 routers.

D. Devices in an IP fabric must be connected to a fabric controller.

**Suggested Answer:** *AC*

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

Question #18 *Topic 2*

Click the Exhibit button.

```
user@spine1# show protocols bgp group underlay
type external;
export Export-Directs;
local-as 65101;
multipath {
    multiple-as;
}
neighbor 172.16.1.1 {
    peer-as 65201;
}
neighbor 172.16.1.5 {
    peer-as 65203;
}
neighbor 172.16.1.3 {
    peer-as 65202;
}
user@spine1# show policy-options
policy-statement Export-Directs {
    term loopback {
        from {
            protocol direct;
            route-filter 192.168.100.0/24 orlonger;
        }
        then accept;
    }
}
```

Referring to the exhibit, the spine1 device has an underlay BGP group that in configured to peer with its neighbors' directly connected interfaces. Which two statements are true in this scenario? (Choose two.)

A. The multihop statement is not required to establish the underlay BGP sessions.

B. Load balancing for the underlay is not configured correctly.

C. Load balancing for the underlay is configured correctly.

D. The multihop statement is required to establish the underlay BGP sessions.

**Suggested Answer:** *AC*

*Community vote distribution*

CD (50%) AC (50%)

**Moscow** 4 weeks, 1 day ago

**Selected Answer: AC**

A - directly connected interfaces

C - multipath multiple-as

upvoted 1 times

**Kornrono** 3 months ago

**Selected Answer: CD**

In the fabric underlay using eBGP, ebgp multihop option needs to be enabled to use loopback addresses for EBGP peering. By default, eBGP assumes directly connected peers, and using loopbacks for peering requires the multihop option. This allows the eBGP sessions to establish and exchange routing information between loopbacks.

upvoted 1 times

Question #19 *Topic 2*

You are using EBGP peering in an underlay IP fabric.
Which two statements are correct in this scenario? (Choose two.)

A. Every leaf node has one peering session to every spine node.

B. Every leaf node has a peering session to every other leaf node

C. EBGP peering does not require an IGP protocol for adjacency establishment.

D. EBGP peering requires an IGP protocol for adjacency establishment.

**Suggested Answer:** *AC*

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

Question #20 *Topic 2*

An IP fabric with EBGP configured in the overlay uses a different AS number for each node in the fabric.
Which two benefits does this configuration have over using a single AS scheme? (Choose two.)

A. There are fewer BGP peers.

B. It avoids the use of route reflectors.

C. There is more TCAM space for AS numbers.

D. It provides more efficient BGP path selection.

**Suggested Answer:** *BD*

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

Click the Exhibit button.

```
user@leaf1> show ethernet-switching vxlan-tunnel-end-point remote
Logical System Name       Id  SVTEP-IP          IFL   L3-Idx     SVTEP-Mode     ELP-
SVTEP-IP
                  0   192.168.100.11   lo0.0     0
 RVTEP-IP         L2-RTT                   IFL-Idx   Interface    NH-Id   RVTEP-Mode
 ELP-IP        Flags
 192.168.100.13   default-switch           571       vtep.32769   1758    RNVE
    VNID          MC-Group-IP
    5010          0.0.0.0
    5020          0.0.0.0
user@leaf1> show interfaces vtep.32769
  Logical interface vtep.32769 (Index 571) (SNMP ifIndex 534)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 192.168.100.13, L2 Routing
Instance: default-switch, L3 Routing Instance: default
    Input packets : 0
    Output packets: 19
  ...
user@leaf1> show evpn database
Instance: default-switch
VLAN  DomainId  MAC address        Active source                  Timestamp
      IP address
     5010       00:00:5e:00:01:01  05:00:00:fd:e9:00:00:13:92:00  Apr 15 22:27:02
 10.1.1.254
     5010       00:0c:29:e8:b7:39  xe-0/0/4.0                     Apr 15 19:41:27
 10.1.1.1
     5010       02:05:86:a7:4c:00  irb.10                         Apr 15 18:50:45
 10.1.1.101
     5020       00:00:5e:00:01:01  05:00:00:fd:e9:00:00:13:9c:00  Apr 15 22:26:51
 10.1.2.254
     5020       00:0c:29:08:04:a0  192.168.100.13                 Apr 15 23:07:22
 10.1.2.1
     5020       02:05:86:a7:4c:00  irb.20                         Apr 15 22:26:51
 10.1.2.101
user@leaf1> show route table bgp.evpn.0 evpn-mac-address 00:0c:29:08:04:a0
bgp.evpn.0: 28 destinations, 42 routes (28 active, 0 holddown, 0 hidden)
Active Route, - = Last Active, * = Both
92.168.100.13:1::5020::00:0c:29:08:04:a0/304 MAC/IP
                *[BGP/170] 00:49:55, localpref 100, from 192.168.100.1
                  AS path: I, validation-state: unverified
                  > to 172.16.1.0 via xe-0/0/0.0
                    to 172.16.1.6 via xe-0/0/1.0
r@leaf1> show route forwarding-table matching 10.1.2.1

tination        Type RtRef Next hop           Type Index    NhRef Netif
1.2.1/32        dest     0 0:c:29:8:4:a0      ucst     1775     1 vtep.32769
```

Referring to the exhibit, Host1 (10.1.1.1) if failing to communicate with Host2 (10.1.2.1) in a data center that uses an ERB architecture. What do you determine from the output?

A. Host1 and Host2 are directly connected to leaf1.

B. The irb.20 interface is not configured on leaf1.

C. The traffic is failing because load balancing is not configured correctly.

D. The traffic is entering the VXLAN tunnel.

**Suggested Answer:** *D*

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

Question #22 *Topic 3*

You are using a single tenant data center with a bridged overlay architecture.
In this scenario, how do hosts of the different virtual networks communicate with each other?

A. using anycast gateway addresses configured on the leaf devices

B. using EVPN Type 5 routes

C. off-fabric using an external device

D. using virtual gateway addresses configured on the spine

**Suggested Answer:** *A*

*Community vote distribution*

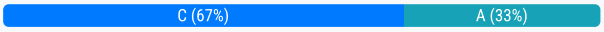


**Moscow** 4 weeks, 1 day ago

**Selected Answer: C**

C. If inter-VLAN routing is required for a bridged overlay, you can use an MX Series router or SRX Series security device that is external to the EVPN/VXLAN fabric.

upvoted 1 times

**basvdw107** 2 months, 2 weeks ago

**Selected Answer: A**

A. using anycast gateway addresses configured on the leaf devices

In a bridged overlay architecture:

Leaf switches act as the first-hop gateways for hosts in different virtual networks
The same IP address (anycast gateway) is configured on multiple leaf switches
Hosts in each virtual network use their local leaf's anycast gateway as their default gateway
Inter-VLAN traffic is routed at the leaf layer where the traffic enters the fabric
This distributed gateway approach eliminates the need for traffic tromboning to spine devices or external routers, providing optimal east-west communication between virtual networks.

The other options don't apply to a bridged overlay architecture:

Option B relates to EVPN-VXLAN routing
Option C would create inefficient traffic patterns
Option D isn't typical as spine devices don't usually provide gateway functionality

upvoted 1 times

**Kornrono** 3 months ago

**Selected Answer: C**

In a bridged overlay, the focus is on providing Layer 2 connectivity between VNIs, rather than relying heavily on Layer 3 routing within the fabric. Spine devices typically don't act as gateways, but rather provide connectivity between leaf devices.

Leaf devices often terminate the VXLAN tunnels and may use external devices for routing between VNIs.

If inter-VLAN routing is required for a bridged overlay, you can use an MX Series router or SRX Series security device that is external to the EVPN/VXLAN fabric.

https://www.juniper.net/documentation/us/en/software/nce/sg-005-data-center-fabric/topics/task/bridged-overlay-cloud-dc-configuring.html#:~:text=A%20bridged%20overlay%20provides%20Ethernet,%2Dto%2Dleaf%20VXLAN%20tunnels.

upvoted 1 times

Question #23 *Topic 3*

What are two ways in which an EVPN-signaled VXLAN is different from a multicast-signaled VXLAN? (Choose two.)

A. An EVPN-signaled VXLAN is less resource intensive.

B. An EVPN-signaled VXLAN can perform autodiscovery of VTEPs using IS-IS.

C. An EVPN-signaled VXLAN can perform autodiscovery of VTEPs using BGP.

D. An EVPN-signaled VXLAN features slower and more complete convergence.

**Suggested Answer:** *CD*

*Community vote distribution*

AC (100%)

**basvdw107** 2 months, 2 weeks ago

**Selected Answer: AC**

Looking at the key differences between EVPN-signaled VXLAN and multicast-signaled VXLAN, the two correct answers are:

A. An EVPN-signaled VXLAN is less resource intensive.
C. An EVPN-signaled VXLAN can perform autodiscovery of VTEPs using BGP.

Let me explain why:
A is correct because EVPN-signaled VXLAN eliminates the need for multicast in the underlay network. Multicast-based VXLAN requires multicast protocols like PIM to be configured throughout the network, which consumes more resources and is more complex to manage.
C is correct because EVPN uses BGP for VTEP discovery and MAC address distribution. BGP EVPN provides a control plane mechanism for VXLAN that enables auto-discovery of VTEPs, distribution of tenant MAC addresses, and establishment of overlay tunnels.
B is incorrect because EVPN uses BGP, not IS-IS, for VTEP autodiscovery.
D is incorrect because EVPN-signaled VXLAN actually provides faster and more complete convergence compared to multicast-signaled VXLAN, not slower.
upvoted 2 times

Question #24 *Topic 3*

Click the Exhibit button.

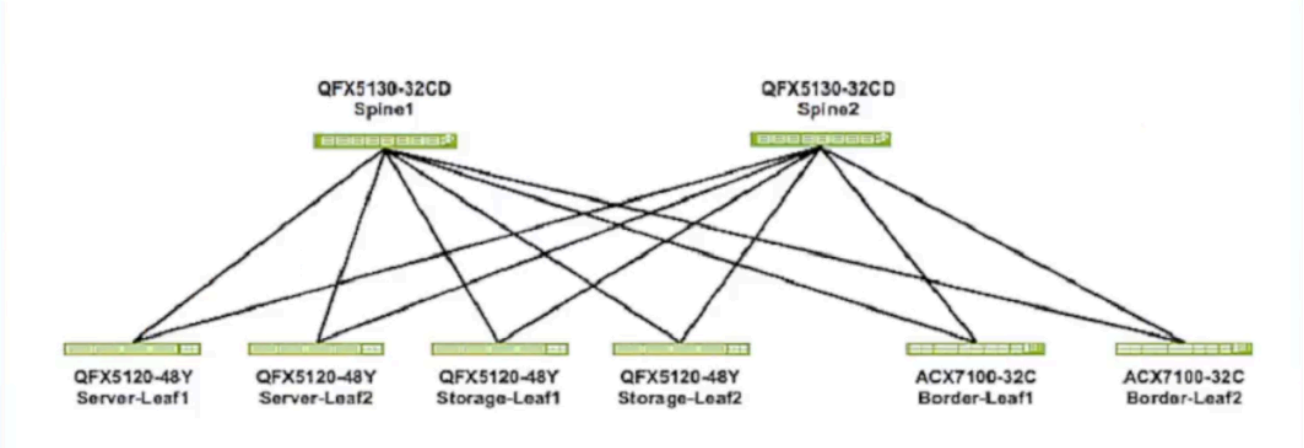


You are deploying a VXLAN overlay with EVPN as the control plane in an ERB architecture.
Referring to the exhibit, which three statements are correct about where then VXLAN gateways will be placed? (Choose three.)

A. Spine devices will have no VXLAN gateways.

B. All leaf devices will have L3 VXLAN gateways.

C. Only the spine devices will have L2 VXLAN gateways.

D. Only the border and leaf devices will have L3 VXLAN gateways.

E. All leaf devices will have L2 VXLAN gateways.

**Suggested Answer:** *ADE*

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

Question #25 *Topic 3*

You are asked to implement VXLAN group-based policies (GBPs) in your data center.
Which two statements are correct in this scenario? (Choose two.)

A. VXLAN GBP ensures consistent application of BGP groups throughout the network.

B. VXLAN GBP ensures consistent application of security group policies throughout the network.

C. VXLAN GBP uses scalable group tags that must be configured statically on each switch and activated through 802.1X.

D. VXLAN GBP uses scalable group tags that may be configured on a RADIUS server and pushed to the switch through 802.1X.

**Suggested Answer:** *BD*

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

Question #26 *Topic 3*

You are asked to configure telemetry on the QFX Series devices in your data center fabric. You want to use sensors that have a vendor-neutral data model.
Which type of sensor should you use in this scenario?

A. JTI native sensors

B. JTI OpenConfig sensors

C. Python sensors

D. analog sensors

**Suggested Answer:** *B*

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

Question #27 *Topic 3*

Click the Exhibit button.

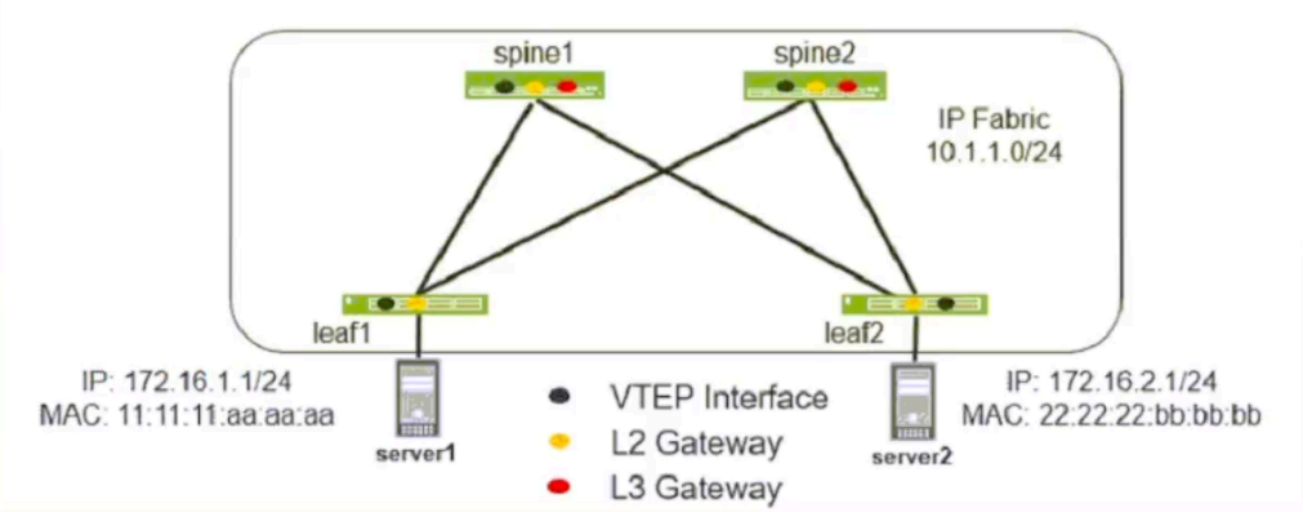


You have implemented an EVPN-VXLAN data center. Device server1 must be able to communicate with device server2.
Referring to the exhibit, which two statements are correct? (Choose two.)

A. Traffic from server1 to server2 will transit a VXLAN tunnel to spine1 or spine2, then a VXLAN tunnel from spine1 or spine2 to leaf2.

B. Traffic from server1 to server2 will transit the VXLAN tunnel between leaf1 and leaf2.

C. An IRB interface must be configured on leaf1 and leaf2.

D. An IRB interface must be configured on spine1 and spine2.

**Suggested Answer:** *AC*

*Community vote distribution*

AD (100%)

**MichelT** 4 months ago

**Selected Answer: AD**

When routing is done in the spines (CRB model) a IREB must be configured at the spines.

upvoted 2 times

Question #28 *Topic 3*

A local VTEP has two ECMP paths to a remote VTEP.
Which two statements are correct when load balancing is enabled in this scenario? (Choose two.)

A. The inner packet fields are not used in the hash load balancing.

B. The destination port in the UDP header is used to load balance VXLAN traffic.

C. The source port in the UDP header is used to load balance VXLAN traffic.

D. The inner packet fields are used in the hash for load balancing.

**Suggested Answer:** *CD*

*Community vote distribution*

CD (100%)

**Moscow** 4 weeks, 1 day ago

**Selected Answer: CD**

https://www.juniper.net/documentation/us/en/software/junos/evpn/topics/concept/evpn-vxlan-dynamic-load-balancing.html
The hashing takes place before a packet undergoes VXLAN encapsulation.

upvoted 1 times

**MichelT** 4 months ago

**Selected Answer: CD**

The source port field in the UDP header is used to enable ECMP load balancing of the VXLAN traffic in the Layer 3 network. This field is set to a hash of the inner packet fields, which results in a variable that ECMP can use to distinguish between tunnels (flows).

upvoted 1 times

Question #29 *Topic 3*

Which three statements are correct about VXLAN control planes? (Choose three.)

A. EVPN is inefficient and does not scale well.

B. EVPN enables fast convergence and updates.

C. Multicast is not agile and requires manual VNI mapping.

D. Multicast does not require as many resources.

E. Both multicast and EVPN can facilitate MAC learning.

**Suggested Answer:** *BDE*

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

Question #30 *Topic 3*

You are implementing VXLAN broadcast domains in your data center environment.
Which two statements are correct in this scenario? (Choose two.)

A. The VNI is a 16-bit value and can range from 0 through 16,777,215.

B. Layer 2 frames are encapsulated by the source VTEP.

C. A VXLAN packet does not contain a VLAN ID.

D. The VNI must match the VLAN tag to ensure that the remote VTEP can decapsulate VXLAN packets.

**Suggested Answer:** *BC*

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

Question #31 *Topic 3*

You want to ensure that VXLAN traffic from the xe-0/0/12 interface is being encapsulated by logical vtep.32770 and sent to a remote leaf device. In this scenario, which command would you use to verify that traffic is flowing?

A. monitor traffic interface xe-0/0/12

B. show interfaces vtep.32770 detail

C. show interface terse vtep.32770

D. show interfaces terse vtep.32770 statistics

**Suggested Answer:** *D*

*Community vote distribution*

B (100%)

**Moscow** 4 weeks, 1 day ago

**Selected Answer: B**

https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/command/show-interfaces-detail.html

upvoted 1 times

Question #32 *Topic 3*

You are asked to build out the VXLAN control plane.
In this scenario, which two protocols provide this capability? (Choose two.)

A. MBGP

B. PIM-SM

C. OSPF

D. IS-IS

**Suggested Answer:** *AD*

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