

DUMPS ARENA

Mist AI Wired Specialist (JNCIS-MistAI-Wired)

Juniper JN0-460

Version Demo

Total Demo Questions: 10

Total Premium Questions: 107

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Wired Assurance	56
Topic 2, Marvis Actions	7
Topic 3, Access Assurance	1
Topic 4, Campus Fabric	42
Topic 5, Mix Questions	1
Total	107

QUESTION NO: 1

You must move from a campus fabric core-distribution centrally-routed bridging (CRB) network to an edge-routed bridging (ERB) network.

In this scenario, where does the gateway for the network move?

- A. from access to distribution
- B. from distribution to core
- C. from distribution to access
- D. from core to distribution

ANSWER: D

Explanation:

In a core-distribution campus fabric, centrally-routed bridging (CRB) places the Layer 3 default gateway, typically implemented by an integrated routing and bridging (IRB) interface, at the core. The distribution devices extend endpoint VLANs toward that centralized gateway using Layer 2 forwarding. Moving to edge-routed bridging (ERB) relocates the gateway function to the edge of the fabric. In this core-distribution design, the distribution layer is the fabric edge because it is the layer directly downstream of the core and closer to the access layer and attached clients. Therefore, the gateway moves **from core to distribution**. This change enables client traffic to be routed at the distribution devices rather than being bridged to the core before routing. It also reduces unnecessary Layer 2 extension and places the first-hop gateway closer to the endpoints, while the routed fabric provides connectivity between distribution devices. Juniper's campus-fabric architecture documentation describes CRB and ERB as designs distinguished by the placement of the Layer 3 gateway, with ERB locating routing at fabric-edge devices. See Juniper's [Campus Fabric documentation](#) and the [EVPN-VXLAN configuration overview](#).

QUESTION NO: 2

What are the two ways to forward BUM traffic when using EVPN?(Choose two.)

- A. Use static routes.
- B. Use BGP neighborship.
- C. Use ingress replication.
- D. Use underlay replication.

ANSWER: C D

Explanation:

Use ingress replication and use underlay replication are the two supported mechanisms for distributing BUM traffic in an EVPN-VXLAN deployment. BUM includes broadcast, unknown-unicast, and multicast Ethernet frames that must be delivered to the relevant remote VTEPs for a VLAN or VNI. With ingress replication, the ingress VTEP makes a separate VXLAN-encapsulated unicast copy of a BUM frame for each remote VTEP that participates in the VNI. EVPN control-plane information identifies those remote VTEPs, allowing replication without requiring multicast routing in the IP fabric. This approach is often useful when the underlay is unicast-only.

With underlay replication, also called multicast replication, the VTEP encapsulates the BUM frame toward an underlay multicast group. The IP underlay replicates the packet along its multicast distribution tree to participating egress VTEPs. This avoids creating a separate copy at the ingress VTEP and can improve replication efficiency at scale, provided the underlay supports the required multicast routing and group membership behavior. Juniper documents both ingress replication and multicast-based BUM delivery for EVPN-VXLAN: [EVPN-VXLAN Overview](#) and [EVPN-VXLAN Multicast](#).

QUESTION NO: 3

Which two steps must be performed when configuring Dynamic Port Configuration (DPC) in the Juniper Mist dashboard? (Choose two.)

- A. Select the Configuration Profile that client devices will be assigned to when matching the DPC.
- B. Define the polling frequency for DPC.
- C. Assign the ports on which you would like to enable Dynamic Configuration.
- D. Define the minimum Junos version.

ANSWER: A C

Explanation:

Dynamic Port Configuration uses an endpoint match to automate the configuration applied to an eligible switch interface. Therefore, selecting the Configuration Profile that client devices will be assigned to when matching the DPC is required: the profile is the reusable definition of the port behavior that DPC applies after a device matches the configured identification criteria. This enables administrators to consistently deliver the intended wired service, such as the appropriate VLAN, authentication behavior, PoE settings, port mode, or other interface parameters, without manually configuring each individual connection.

Administrators must also assign the ports on which they would like to enable Dynamic Configuration. DPC is deliberately scoped to selected switch interfaces so that automation is applied only where endpoint-driven configuration is desired. The selected interfaces act as the set of ports on which Mist evaluates connected-device information and applies the associated Configuration Profile when a match occurs. Together, the profile selection and port assignment define both the configuration outcome and the interfaces where that outcome can be automatically enforced. Juniper's Mist wired documentation describes switch-port configuration and profile-based policy management at [Juniper Mist Wired Port Configuration](#); related Mist Wired documentation is available at [Juniper Mist Wired Documentation](#).

QUESTION NO: 4

What are three wired Service Level Expectations (SLEs) that comprise Juniper Mist Wired Assurance? (Choose three.)

- A. Capacity
- B. Successful Connects
- C. Switch Health
- D. Coverage
- E. Throughput

ANSWER: B C E

Explanation:

The wired SLEs in Juniper Mist Wired Assurance are **Successful Connects**, **Switch Health**, and **Throughput**. These measurements provide an experience-focused view of the wired environment, allowing administrators to identify degraded service and investigate the contributing network conditions from the Mist dashboard. Successful Connects measures whether endpoints can successfully establish wired connectivity, helping reveal connection-related user-impacting events. Switch Health provides visibility into the operational condition of managed switches, including factors that can affect reliable service delivery. Throughput measures the quality of traffic delivery across the wired network and helps identify performance degradation affecting connected clients and applications. Together, these SLEs turn raw switch and client telemetry into actionable service-level insight, supporting quicker detection, correlation, and remediation of wired user-experience issues. Juniper describes Wired Assurance as an AI-native operations capability that uses SLEs and event correlation to simplify wired troubleshooting and provide visibility into network performance. See the [Juniper Mist Wired Assurance documentation](#) and the [Juniper Mist Wired Assurance product page](#).

QUESTION NO: 5

Which statement is correct about a 3-stage campus fabric IP Clos?

- A. The distribution layer is connected to the access layer.
- B. The distribution layer is connected to the core layer.
- C. The core layer is connected to the access layer.
- D. The core layer devices are connected to each other.

ANSWER: C

Explanation:

The core layer is connected to the access layer. In Juniper Mist Wired Assurance campus-fabric terminology, a 3-stage IP Clos design is intended for smaller deployments that do not require a distribution layer. The fabric therefore provides direct routed connectivity between the access switches and the core switches. This is a Clos-based leaf-and-spine-style design: access-layer devices form the edge of the fabric and core devices provide the transit layer between access devices as needed.

This direct core-to-access relationship is central to the simplified 3-stage deployment model. Mist automates the fabric configuration, including the routed underlay and the Layer 3 gateway configuration at the access layer. Access switches use integrated routing and bridging interfaces for attached VLANs, while the interswitch fabric links provide IP connectivity toward the core. The design supports scalable, predictable forwarding while avoiding the additional intermediate distribution tier that is used in larger campus-fabric designs.

Juniper describes Campus Fabric as an IP Clos architecture managed through Mist Wired Assurance. See the [Juniper Campus Fabric documentation](#) and the [Juniper Mist AI product overview](#).

QUESTION NO: 6

You want a Juniper cloud-ready switch to connect to the Juniper Mist Cloud.

In this scenario, which two statements are correct?(Choose two.)

- A. The switch must be running Junos OS version 23.2 or higher.
- B. The switch needs outbound connectivity to the Juniper Mist Cloud using either TCP port 2200 or port 443.
- C. The switch must connect to the Juniper Agile Licensing Portal.
- D. The switch must connect to a DNS server.

ANSWER: B D

Explanation:

The switch needs outbound connectivity to the Juniper Mist Cloud using either TCP port 2200 or port 443 because a cloud-managed EX Series switch initiates its management connection outbound to the Mist cloud. This outbound session enables the switch to register with the assigned organization, retrieve its configuration, maintain its cloud-management connection, and send operational telemetry. The firewall or upstream security policy must therefore permit the applicable outbound TCP communication from the switch management network; no inbound connection initiated from the Internet to the switch is required for this cloud-management relationship.

The switch must connect to a DNS server because it must resolve the Mist cloud service hostnames before it can establish its outbound cloud connection. DNS reachability is consequently a foundational onboarding dependency, alongside IP connectivity, a default route where required, and the permitted outbound service ports. In practical deployments, administrators commonly provide DNS through DHCP or static configuration on the switch management interface, then verify name resolution and Internet or cloud-path reachability before claiming or onboarding the device. Juniper's cloud-ready switching documentation describes the required network connectivity for cloud-managed switches, and the Mist Wired

Assurance documentation provides the operational context for managing EX switches from the cloud. See [Juniper cloud-ready switches documentation](#) and [Mist Wired Assurance Administration Guide](#).

QUESTION NO: 7

Which statement about RadSec is correct?

- A. RadSec allows RADIUS authentication over TCP and TLS.
- B. RadSec allows RADIUS authentication over TCP and SSL.
- C. RadSec allows RADIUS authentication over UDP and TLS.
- D. RadSec allows RADIUS authentication over UDP and SSL.

ANSWER: A

Explanation:

RadSec allows RADIUS authentication over TCP and TLS. RadSec is the standardized method for transporting RADIUS traffic through a TLS-protected TCP connection. Traditional RADIUS commonly uses UDP, and although it includes limited protection for certain attributes, it does not provide the connection-wide confidentiality, strong peer authentication, and integrity protection available with TLS. By placing RADIUS protocol exchanges inside TLS, RadSec protects authentication and accounting communication between a network access device and a RADIUS server against interception and modification while also allowing certificate-based authentication of the communicating endpoints.

This is particularly useful where RADIUS traffic must traverse untrusted or routed networks, or where organizations require stronger protection for wired-access authentication services such as IEEE 802.1X. TCP provides the reliable, connection-oriented transport used by the RadSec specification, while TLS establishes the cryptographic security context over that transport. The IETF defines this protocol in RFC 6614, including its use of TCP and TLS and the default RadSec port. Juniper's documentation also describes RadSec as RADIUS over TLS for securing communication with RADIUS infrastructure. See [RFC 6614: Transport Layer Security \(TLS\) Encryption for RADIUS](#) and [Juniper RADIUS Authentication documentation](#).

QUESTION NO: 8

You are designing a core-distribution campus fabric with Layer 2-only access switches. Most traffic is north-south traffic that must be routed toward shared services connected at the core tier.

Which gateway-placement model is most appropriate in this scenario?

- A. campus fabric core-distribution with centrally-routed bridging (CRB)
- B. campus fabric core-distribution with edge-routed bridging (ERB)
- C. campus EVPN multihoming
- D. campus fabric IP Clos

ANSWER: A

Explanation:

Centrally-routed bridging (CRB) is correct because the integrated routing and bridging gateway is placed centrally in the fabric, typically at the core tier. Traffic requiring inter-VLAN or inter-subnet routing is therefore forwarded toward the central gateway.

Edge-routed bridging places the gateway at the distribution edge and is better suited to designs that benefit from routing east-west traffic close to the access layer. EVPN multihoming provides attachment redundancy, but it does not define the placement of the Layer 3 gateway. An IP Clos design extends the fabric to the access tier rather than retaining a conventional Layer 2 access tier.

QUESTION NO: 9

You are asked to apply the same system-level configuration across all the devices in multiple sites using Mist AI.

According to Juniper Networks, which solution should you use in this scenario?

- A. Use the CLI on each device.
- B. Use the site-level switch configuration option in Mist AI.
- C. Use an organization-level template in Mist AI.
- D. Use the individual switch configuration option in Mist AI.

ANSWER: C

Explanation:

Use an organization-level template in Mist AI. Mist's wired configuration model lets administrators define reusable switch configuration templates at the organization scope and assign them to multiple sites. This provides a centralized, repeatable way to distribute common system-level settings consistently wherever the template is applied. A template can include switch configuration elements such as system services and other shared configuration required by the organization, while still allowing site-specific and device-specific settings to be managed at the appropriate lower scopes when necessary.

This approach is especially appropriate when the objective is configuration standardization across more than one site. Rather than reproducing the same settings separately for every location, the organization-level template establishes a common baseline and allows Mist to manage deployment to the switches associated with the selected sites. It also improves operational consistency: updates to the shared baseline can be made centrally and propagated through the template assignment workflow. Juniper's Mist wired documentation describes configuration templates as organization-level objects that can be assigned to sites, and its Wired Assurance documentation explains the centralized cloud-based configuration and management model for EX switches. See [Juniper Mist Wired switch configuration documentation](#) and [Juniper Wired Assurance overview](#).

QUESTION NO: 10

You receive a Marvis Action that reports "Port negotiation mismatch detected on EX3400-48P."

In this scenario, which two steps should you take to resolve the issue? (Choose two.)

- A. Verify that both switch ports are set to auto-negotiate speed and duplex.
- B. Force the connected device port to half-duplex mode.
- C. Replace the switch uplink SFP module.
- D. Check that the connected device supports 802.3 autonegotiation.

ANSWER: A D

Explanation:

A port-negotiation mismatch means the two ends of an Ethernet link have not established compatible speed and duplex parameters. The appropriate first step is to verify that both switch ports are set to auto-negotiate speed and duplex. Autonegotiation allows the local EX3400 interface and its peer to advertise their supported capabilities and select a mutually supported operational link mode. This avoids duplex conflicts and speed mismatches that can produce errors, poor throughput, or an unusable link.

It is also necessary to check that the connected device supports 802.3 autonegotiation. Both link partners must be capable of, and configured consistently for, autonegotiation for the normal negotiation process to succeed. Once support and configuration are confirmed on both ends, the interfaces can renegotiate a common speed and duplex setting. Juniper documents Ethernet autonegotiation behavior and configuration considerations for EX and QFX platforms in its support

guidance: [Understanding Auto-Negotiation on EX and QFX Switches](#). The relevant standard is IEEE 802.3 Ethernet, whose working-group resources are available at [IEEE 802.3 Ethernet Working Group](#).