

Cisco 300-615


Troubleshooting Cisco Data Center Infrastructure

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/300-615>

Latest Version: 20.0

Question: 1

A vPC Type-1 inconsistency between two vPC peers in a VXLAN EVPN setup is discovered. Which two actions need to be attempted to resolve the issue? (Choose two.)

- A. Configure the NVE interfaces to be Up on both switches.
- B. Set a different distributed gateway virtual MAC address.
- C. Set a different secondary IP addresses on NVE source-interface.
- D. Configure the same VNI to multicast group mapping.
- E. Set a different primary IP addresses on NVE source-interface.

Answer: A, D

Explanation:

Configure the NVE interfaces to be Up on both switches; Configure the same VNI to multicast group mapping is the valid corrective action because the failure described is a protocol or fabric state problem, not a generic connectivity issue. For vPC, VXLAN, EVPN, NVE, VNI, Cisco troubleshooting expects the engineer to compare both endpoints of the relationship: vPC peers must agree on consistency parameters, HSRP peers must exchange hellos in the same VLAN and version context, OSPF/BGP peers must match the attributes required for stable adjacency, and VXLAN EVPN must advertise the correct host or prefix information. The selected option restores that required state. The other options either operate on the wrong layer, change a value that is not used to form the relationship, or would make the peer relationship less deterministic. A clean validation is to inspect the control-plane table first, then verify the forwarding table, and finally test end-to-end traffic. This is the Cisco-preferred approach because it avoids random changes and confirms that the fabric is forwarding based on correct protocol state.

Question: 2

A mission-critical server is connected to site A. Connectivity to this server is lost from site B because the MAC route is missing in the OTV VDC of the Nexus 7000 in site B due to MAC aging. Which action allows the flooding of the unknown unicast MAC on the Nexus 7000 in the OTV VDC?

- A. Use route-map to advertise this MAC statically and redistribute with ISIS.
- B. Unknown unicast flooding is not allowed.
- C. Use the `otv flood mac <>` command to selectively flood traffic for a given MAC.
- D. Use the `otv isis bfd <>` command to configure BFD protocol.

Answer: C

Explanation:

The answer is Use the `otv flood mac <>` command to selectively flood traffic for a given MAC because Cisco Nexus and ACI forwarding depends on exact agreement between the devices participating in the feature. The scenario references OTV; those technologies fail in predictable ways when required parameters do not match or when the required control-plane advertisement is absent. In DCIT troubleshooting, the engineer should identify the failed dependency, correct that dependency, and then verify the operational table that proves the fix. For example, vPC and EVPN checks focus on consistency, VNI/VRF bindings, peer-gateway behavior, and host-route advertisement; routing checks focus on neighbor state, MTU, timers, authentication, next-hop, and route preference. The selected answer addresses the dependency that blocks forwarding. The distractors are plausible because they touch related features, but they do not resolve the specific failure condition. Post-change validation should include protocol state, learned endpoint or route entries, and a traffic test across the affected path.

Question: 3

Refer to the exhibit.

```
Switch-B# show vpc
vPC domain id: 34
Peer status : peer adjacency formed ok
vPC keep-alive status : peer is alive
Configuration consistency status : success
Per-vlan consistency status : success
Type-2 consistency status : success
vPC role : primary, operational secondary
Number of vPCs configured : 49
Peer Gateway : Enabled
Peer gateway excluded VLANs: -
Dual-active excluded VLANs : -
Graceful Consistency Check: Enabled
Auto-recovery status: Enabled (timeout = 240 seconds)
```

After a failover occurs, which two actions must be performed on Switch-B to manually preempt the operational primary role back to Switch-A? (Choose two.)

- A. Configure the local vPC role priority to have a lower value than Switch-A.
- B. Configure the local vPC role priority to have a higher value than Switch-A.
- C. Disable and then re-enable the vPC peer-keepalive link.
- D. Configure the local vPC role priority to have the same value as Switch-A.
- E. Disable and then re-enable the vPC peer link.

Answer: B, E

Explanation:

Cisco troubleshooting logic supports Configure the local vPC role priority to have a higher value than Switch-A; Disable and then re-enable the vPC peer link as the required fix. In this question the symptoms point to vPC, and the defect is best handled by correcting the exact control-plane or consistency condition. Nexus platforms can show interfaces or peers as configured while still refusing to forward traffic if a vPC Type-1 parameter, VLAN allowance, MTU, HSRP version, EVPN route, or routing attribute is wrong. Cisco methodology is to avoid broad resets unless the state machine is stuck; instead, correct the mismatched parameter and verify that the resulting operational state becomes stable. The other choices do not remove the root cause, or they modify a setting that is not used by the failing protocol exchange. After the change, the engineer should review the peer status, route or MAC table, and counters for drops or inconsistency messages. That sequence proves that the network fabric is fixed, not merely restarted.

Question: 4

Refer to the exhibit.

```
Debug messages from Router-A
OSPF: Rcv DBD from 10.100.1.2 on GigabitEthernet0/1 seq 0x2124 opt 0x52 flag 0x2
      len 1452 mtu 2000 state EXSTART
OSPF: Nbr 10.100.1.2 has larger interface MTU
SPF: Send DBD to 10.100.1.2 on GigabitEthernet0/1 seq 0x9E6 opt 0x52 flag 0x7
      len 32
OSPF: Retransmitting DBD to 10.100.1.2 on GigabitEthernet0/1 [10]
OSPF: Send DBD to 10.100.1.2 on GigabitEthernet0/1 seq 0x9E6 opt 0x52 flag 0x7
      len 32
OSPF Retransmitting DBD to 10.100.1.2 on GigabitEthernet0/1[11]
%OSPF-5-ADJCHG: Process 1, Nbr 10.100.1.2 on GigabitEthernet0/1 from EXSTART to
      DOWN, Neighbor Down: Too many retransmissions

Debug messages from Router-B
OSPF: Rev DBD from 10.100.100.1 on GigabitEthernet0/1 seq 0x89E opt 0x52 flag 0x7
      len 32 mtu 1600 state EXCHANGE
OSPF: Nbr 10.100.100.1 has smaller interface MTU
```

An OSPF adjacency between Router-A and Router-B cannot reach the FULL state. Which action resolves the issue?

- A. Adjust the MTU on Router-A to 1600.
- B. Disable the check of the MTU value.
- C. Set the OSPF media type to point-to-point.
- D. Adjust the MTU on Router-B to 1604.

Answer: B

Explanation:

The correct remediation is Disable the check of the MTU value. In Cisco data center troubleshooting, the first priority is to isolate whether the failure is caused by the underlay, overlay control plane, first-hop gateway behavior, or a consistency parameter. This item is centered on OSPF, so the selected action targets the condition that prevents normal forwarding rather than masking the symptom. Cisco Nexus and ACI guidance treats vPC, VXLAN EVPN, OTV, OSPF, BGP, HSRP, and LACP issues as state-dependent technologies: a neighbor, peer, VLAN, VNI, multicast group, timer, MTU, or policy mismatch can leave the protocol partially operational while traffic is still dropped. The rejected alternatives either change unrelated attributes, introduce a new inconsistency, or affect only a secondary behavior. The operational check after applying the fix should be to confirm adjacency or peer state, validate the relevant route/MAC/endpoint table, and test traffic in the affected VLAN or VRF. This aligns with Cisco DCIT network troubleshooting methodology for isolating control-plane state before changing data-plane policy.

Question: 5

The Cisco Nexus switch is connected to a peer switch that is not running Cisco NX-OS. The switches are connected using port channel and are experiencing packet loss. Which action should be performed on the Cisco Nexus switch ports to resolve this issue?

- A. Turn on lacp suspend-individual.
- B. Turn on lacp graceful-convergence.
- C. Turn off lacp graceful-convergence.
- D. Turn off lacp suspend-individual.

Answer: C

Explanation:

Turn off lacp graceful-convergence is the valid corrective action because the failure described is a protocol or fabric state problem, not a generic connectivity issue. For LACP, Cisco Nexus, Cisco troubleshooting expects the engineer to compare both endpoints of the relationship: vPC peers must agree on consistency parameters, HSRP peers must exchange hellos in the same VLAN and version context, OSPF/BGP peers must match the attributes required for stable adjacency, and VXLAN EVPN must advertise the correct host or prefix information. The selected option restores that required state. The other options either operate on the wrong layer, change a value that is not used to form the relationship, or would make the peer relationship less deterministic. A clean validation is to inspect the control-plane table first, then verify the forwarding table, and finally test end-to-end traffic. This is the Cisco-preferred approach because it avoids random changes and confirms that the fabric is forwarding based on correct protocol state.

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>