

## [2017-July-NewExam Pass 100%!Braindump2go 642-997 Exam PDF 137q Instant Download[21-30

2017 July CCNP 642-997 Exam Dumps with PDF and VCE New Updated in [www.Braindump2go.com](#) **Today!100% Real Exam Questions! 100% Exam Pass Guaranteed!**]

1.[2017 New CCNP 642-997 Exam Dumps (PDF & VCE) 137Q&As Download:

<http://www.braindump2go.com/642-997.html> 2.[2017 New CCNP 642-997 Exam Questions & Answers:

<https://drive.google.com/drive/folders/0B75b5xYLjSSNTDVuYIJWQVZ3RkU?usp=sharing> QUESTION 21 Which statement about electronic programmable logic device image upgrades is true? A. EPLD and ISSU image upgrades are nondisruptive. B. An EPLD upgrade must be performed during an ISSU system or kickstart upgrade. C. Whether the module being upgraded is online or offline, only the EPLD images that have different current and new versions are upgraded. D. You can execute an upgrade or downgrade only from the active supervisor module. Answer: D Explanation: You can upgrade (or downgrade) EPLDs using CLI commands on the Nexus 7000 Series device.

[http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/4\\_0/epld/release/notes/epld\\_rn.html](http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/4_0/epld/release/notes/epld_rn.html) QUESTION 22 Which statement about SNMP support on Cisco Nexus switches is true? A. Cisco NX-OS only supports SNMP over IPv4. B. Cisco NX-OS supports one instance of the SNMP per VDC. C. SNMP is not VRF-aware. D. SNMP requires the

LAN\_ENTERPRISE\_SERVICES\_PKG license. E. Only users belonging to the network operator RBAC role can assign SNMP groups. Answer: B Explanation: Cisco NX-OS supports one instance of the SNMP per virtual device context (VDC). By default, Cisco NX-OS places you in the default VDC. SNMP supports multiple MIB module instances and maps them to logical network entities. SNMP is also VRF aware. You can configure SNMP to use a particular VRF to reach the SNMP notification host receiver. You can also configure SNMP to filter notifications to an SNMP host receiver based on the VRF where the notification occurred.

[http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5\\_x/nx-os/system\\_management/configuration/guide/sm\\_nx\\_os\\_cg/sm\\_9snmp.html](http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/system_management/configuration/guide/sm_nx_os_cg/sm_9snmp.html) QUESTION 23 Which GLBP load-balancing algorithm ensures that a client is always mapped to the same VMAC

address? A. vmac-weighted B. dedicated-vmac-mode C. shortest-path and weighting D. host-dependent Answer: D Explanation: Host dependent--GLBP uses the MAC address of the host to determine which virtual MAC address to direct the host to use. This algorithm guarantees that a host gets the same virtual MAC address if the number of virtual forwarders does not change.

[http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5\\_x/nx-os/unicast/configuration/guide/l3\\_cli\\_nxos/l3\\_glb.html](http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/unicast/configuration/guide/l3_cli_nxos/l3_glb.html)

QUESTION 24 What is the grace period in a graceful restart situation? A. how long the supervisor waits for NSF replies B. how often graceful restart messages are sent after a switchover C. how long NSF-aware neighbors should wait after a graceful restart has started before tearing down adjacencies D. how long the NSF-capable switches should wait after detecting that a graceful restart has started, before verifying that adjacencies are still valid Answer: C Explanation: Graceful restart (GR) refers to the capability of the control plane to delay advertising the absence of a peer (going through control-plane switchover) for a "grace period," and thus help minimize disruption during that time (assuming the standby control plane comes up). GR is based on extensions per routing protocol, which are interoperable across vendors. The downside of the grace period is huge when the peer completely fails and never comes up, because that slows down the overall network convergence, which brings us to the final concept: nonstop routing (NSR). NSR is an internal (vendor-specific) mechanism to extend the awareness of routing to the standby routing plane so that in case of failover, the newly active routing plane can take charge of the already established sessions.

<http://www.ciscopress.com/articles/article.asp?p=1395746&seqNum=2> QUESTION 25 Which two types of traffic are carried over a vPC peer link when no failure scenarios are present? (Choose two.) A. multicast data traffic B. unicast data traffic C. broadcast data traffic D. vPC keep-alive messages Answer: A Explanation: The vPC peer link is the link used to synchronize states between the vPC peer devices. The vPC peer link carries control traffic between two vPC switches and also multicast, broadcast data traffic. In some link failure scenarios, it also carries unicast traffic. You should have at least two 10 Gigabit Ethernet interfaces for peer links. [http://www.cisco.com/c/en/us/products/collateral/switches/nexus-5000-series-switches/configuration\\_guide\\_c07-543563.html](http://www.cisco.com/c/en/us/products/collateral/switches/nexus-5000-series-switches/configuration_guide_c07-543563.html)

QUESTION 26 A Cisco Nexus 2000 Series Fabric Extender is connected to two Cisco Nexus 5000 Series switches via a vPC link. After both Cisco Nexus 5000 Series switches lose power, only one switch is able to power back up. At this time, the Cisco Nexus 2000 Series Fabric Extender is not active and the vPC ports are unavailable to the network. Which action will get the Cisco Nexus 2000 Series Fabric Extender active when only one Cisco Nexus 5000 Series switch is up and active? A. Move the line from the failed Cisco Nexus 5000 Series switch to the switch that is powered on, so the port channel forms automatically on the switch that is powered on. B. Shut down the peer link on the Cisco Nexus 5000 Series switch that is powered on. C. Configure reload restore or auto-recovery reload-delay on the Cisco Nexus 5000 Series switch that is powered on. D. Power off and on the Cisco Nexus 2000 Series Fabric Extender so that it can detect only one Cisco Nexus 5000 Series switch at power up. Answer: C Explanation: The vPC

consistency check message is sent by the vPC peer link. The vPC consistency check cannot be performed when the peer link is lost. When the vPC peer link is lost, the operational secondary switch suspends all of its vPC member ports while the vPC member ports remain on the operational primary switch. If the vPC member ports on the primary switch flaps afterwards (for example, when the switch or server that connects to the vPC primary switch is reloaded), the ports remain down due to the vPC consistency check and you cannot add or bring up more vPCs. Beginning with Cisco NX-OS Release 5.0(2)N2(1), the auto-recovery feature brings up the vPC links when one peer is down.

[http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/operations/n5k\\_vpc\\_ops.html](http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/operations/n5k_vpc_ops.html) QUESTION 27 Which policy-map action performs congestion avoidance? A. priority B. bandwidth C. queue-limit D. random-detect Answer: D Explanation: Congestion avoidance techniques monitor network traffic loads in an effort to anticipate and avoid congestion at common network bottlenecks. Congestion avoidance is achieved through packet dropping. Among the more commonly used congestion avoidance mechanisms is Random Early Detection (RED), which is optimum for high-speed transit networks. Cisco IOS QoS includes an implementation of RED that, when configured, controls when the router drops packets. If you do not configure Weighted Random Early Detection (WRED), the router uses the cruder default packet drop mechanism called tail drop.

[http://www.cisco.com/c/en/us/td/docs/ios/12\\_2/qos/configuration/guide/fqos\\_c/qcfcnav.html](http://www.cisco.com/c/en/us/td/docs/ios/12_2/qos/configuration/guide/fqos_c/qcfcnav.html) QUESTION 28 Refer to the exhibit. Which statement based on these two outputs that were collected 24 hours apart is true? A. The Site 2 OTV edge device has gone down. B. The MAC address cannot be discovered on two separate port channel interfaces. C. The MAC address that ends in 020a moved to the local site 23 hours ago. D. The Overlay1 IP address should be a multicast IP address. Answer: C QUESTION 29 Which two reasons explain why a server on VLAN 10 is unable to join a multicast stream that originates on VLAN 20? (Choose two.) A. IGMP snooping and mrouter are not enabled on VLAN 10. B. VLAN 20 has no IGMP snooping querier defined and VLAN 10 has no mrouter. C. The mrouter on VLAN 20 does not see the PIM join. D. The mrouter must be on VLAN 10 and VLAN 20. Answer: AC Explanation: IGMP snooping is a mechanism to constrain multicast traffic to only the ports that have receivers attached. The mechanism adds efficiency because it enables a Layer 2 switch to selectively send out multicast packets on only the ports that need them. Without IGMP snooping, the switch floods the packets on every port. The switch "listens" for the exchange of IGMP messages by the router and the end hosts. In this way, the switch builds an IGMP snooping table that has a list of all the ports that have requested a particular multicast group. The mrouter port is simply the port from the switch point of view that connects to a multicast router. The presence of at least one mrouter port is absolutely essential for the IGMP snooping operation to work across switches. All Catalyst platforms have the ability to dynamically learn about the mrouter port. The switches passively listen to either the Protocol Independent Multicast (PIM) hellos or the IGMP query messages that a multicast router sends out periodically. <http://www.cisco.com/c/en/us/support/docs/switches/catalyst-6500-series-switches/68131-cat-multicast-prob.html>

QUESTION 30 Which two issues explain why a packet is not being routed as desired in a policy-based routing configuration? (Choose two.) A. The route map is not applied to the egress interface. B. The route map is not applied to the ingress interface. C. The next hop that is configured in the route map is not in the global routing table. D. The next hop that is configured in the route map has a higher metric than the default next hop. Answer: CD Explanation: The next hop that is configured in the route map is not in the global routing table then the packet will not be forwarded as desired. The next hop that is configured in the route map has a higher metric than the default next hop. !!!RECOMMEND!!! 1. | 2017 New CCNP 642-997 Exam Dumps (PDF & VCE)

137 Q&As Download: <http://www.braindump2go.com/642-997.html> 2. | 2017 New CCNP 642-997 Study Guide Video: YouTube Video: [YouTube.com/watch?v=P-9rGHWsmU8](https://www.youtube.com/watch?v=P-9rGHWsmU8)