642-997 - DCUFI Implementing Cisco Data Center Unified Fabric (DCUFI) v5.0

http://www.certleader.com/642-997-dumps.html

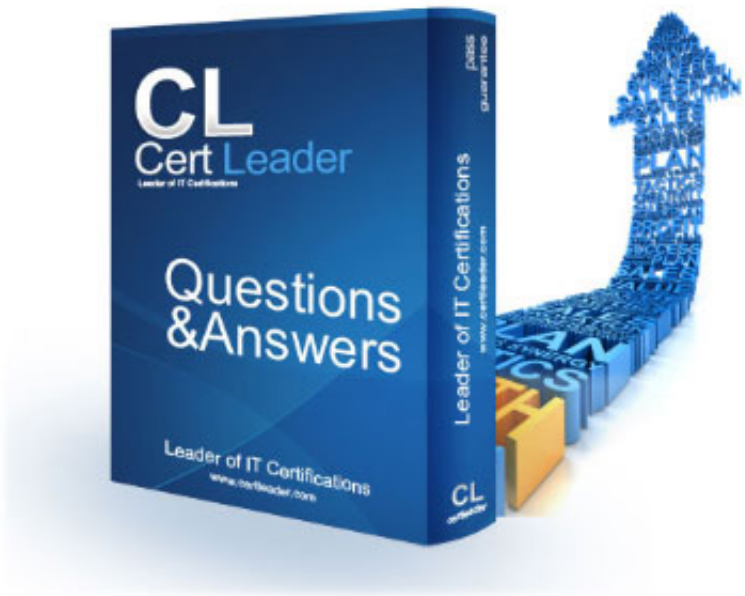

1.Which SCSI terminology is used to describe source and destination nodes?

A. hosts and targets

B. initiators and targets

C. HBA and disks

D. initiators and disks

E. HBA and targets

Answer: B

Explanation:

In computer data storage, a SCSI initiator is the endpoint that initiates a SCSI session, that is, sends a SCSI command. The initiator usually does not provide any Logical Unit Numbers (LUNs).

On the other hand, a SCSI target is the endpoint that does not initiate sessions, but instead waits for initiators' commands and provides required input/output data transfers. The target usually provides to the initiators one or more LUNs, because otherwise no read or write command would be possible.

Reference: http://en.wikipedia.org/wiki/SCSI_initiator_and_target

2. In policy-based routing, which action is taken for packets that do not match any of the route-map statements?

A. forwarded after the egress queue empties on the outbound interface

B. forwarded using the last statement in the route map

C. forwarded using the closest matching route-map statement

D. forwarded using destination-based routing

Answer: D

Explanation:

Each entry in a route map contains a combination of match and set statements. The match statements define the criteria for whether appropriate packets meet the particular policy (that is, the conditions to be met). The set clauses explain how the packets should be routed once they have met the match criteria.

You can mark the route-map statements as permit or deny. You can interpret the statements as follows:

• If the statement is marked as permit and the packets meet the match criteria, the set clause is applied. One of these actions involves choosing the next hop.

• If a statement is marked as deny, the packets that meet the match criteria are sent back through the normal forwarding channels, and destination-based routing is performed.

• If the statement is marked as permit and the packets do not match any route-map statements, the packets are sent back through the normal forwarding channels, and destination-based routing is performed.

Reference: http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/unicast/configuration/guide/l3_cli_nxos/l3pbr.pdf

3. Refer to the exhibit.

```
N7K-1(config)# feature vpc
N7K-1(config)# vpc domain 113
N7K-1(config-vpc-domain)# peer-gateway
N7K-1(config-vpc-domain)#

N7K-2(config)# feature vpc
N7K-2(config)# vpc domain 113
N7K-2(config-vpc-domain)# peer-gateway
N7K-2(config-vpc-domain)#
```

What is the consequence of configuring peer-gateway on the two vPC peers N7K-1 and N7K-2?

A. Nothing, this is the standard vPC configuration to make the feature work.

B. The downstream device detects only one of the vPC peers as its gateway.

C. The downstream device can use DMAC of N7K-1 on the link to N7K-2, and N7K-2 forwards the packet.

D. This configuration enables the downstream device to use DHCP to obtain its default gateway.

Answer:: C

Explanation:

Beginning with Cisco NX-OS 4.2(1), you can configure vPC peer devices to act as the gateway even for packets that are destined to the vPC peer device's MAC address. Use the peer-gateway command to configure this feature.

Some network-attached storage (NAS) devices or load-balancers may have features aimed to optimize the performances of particular applications. Essentially these features avoid performing a routing-table lookup when responding to a request that originated form a host not locally attached to the same subnet. Such devices may reply to traffic using the MAC address of the sender Cisco Nexus 7000 device rather than the common HSRP gateway. Such behavior is non-complaint with some basic Ethernet RFC standards. Packets reaching a vPC device for the non-local router MAC address are sent across the peer-link and could be dropped by the built in vPC loop avoidance mechanism if the final destination is behind another vPC.

The vPC peer-gateway capability allows a vPC switch to act as the active gateway for packets that are addressed to the router MAC address of the vPC peer. This feature enables local forwarding of such packets without the need to cross the vPC peer-link. In this scenario, the feature optimizes use of the peer-link and

avoids potential traffic loss. Configuring the peer-gateway feature needs to be done on both primary and secondary vPC peers and is non-disruptive to the operations of the device or to the vPC traffic. The vPC peer-gateway feature can be configured globally under the vPC domain submode. When enabling this feature it is also required to disable IP redirects on all interface VLANs mapped over a vPC VLAN to avoid generation of IP redirect messages for packets switched through the peer gateway router. When the feature is enabled in the vPC domain, the user is notified of such a requirement through an appropriate message.

Packets arriving at the peer-gateway vPC device will have their TTL decremented, so packets carrying TTL = 1 may be dropped in transit due to TTL expire. This needs to be taken into account when the peer-gateway feature is enabled and particular network protocols sourcing packets with TTL = 1 operate on a vPC VLAN.

Reference: http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/4_2/nx-os/interfaces/configuration/guide/if_nxos/if_vPC.html

4. 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,C

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.

Reference: http://www.cisco.com/c/en/us/products/collateral/switches/nexus-5000-series-switches/configuration_guide_c07-543563.html

5. 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. This feature performs two operations:

•If both switches reload, and only one switch boots up, auto-recovery allows that switch to assume the role of the primary switch. The vPC links come up after a configurable period of time if the vPC peer-link and the peer-keepalive fail to become operational within that time. If the peer-link comes up but the peer-keepalive does not come up, both peer switches keep the vPC links down. This feature is similar to the reload restore feature in Cisco NX- OS Release 5.0(2)N1(1) and earlier releases. The reload delay period can range from 240 to 3600 seconds.

•When you disable vPCs on a secondary vPC switch because of a peer-link failure and then the primary vPC switch fails, the secondary switch reenables the vPCs. In this scenario, the vPC waits for three consecutive keepalive failures before recovering the vPC links.

Reference: http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/operations/n5k_vpc_ops.html

6. 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.

Reference: http://www.ciscopress.com/articles/article.asp?p=1395746&seqNum=2

7. Refer to the exhibit.

```
N7K-1#show fabricpath switch id
FABRICPATH SWITCH-ID TABLE
Legend: '*' - this system
==========================================================================
SWITCH-ID SYSTEM-ID      FLAGS  STATE  STATIC EMULATED
-----------------------------------+----------+-----------+----------+-------------------
  1     0022.5579.b1c1 Primary Confirmed Yes   No
  2     0022.5579.b1c2 Primary Confirmed Yes   No
  3     001b.54c2.7f41 Primary Confirmed Yes   No
  4     001b.54c2.7f42 Primary Confirmed Yes   No
  5     0005.73b1.f0c1 Primary Confirmed Yes   No
 *6     0005.73af.08bc Primary Confirmed Yes   No
  7     0005.73b2.0fbc Primary Confirmed Yes   No
  8     0005.73af.0ebc Primary Confirmed Yes   No
 102     0005.73af.0ebc Primary Confirmed No    Yes
 101     0005.73b2.0fbc Primary Confirmed No    Yes
```

Which three statements about the Cisco Nexus 7000 switch are true? (Choose three.)

A. An emulated switch ID must be unique when the vPC+ feature is used.

B. Switches with FabricPath and vPC+ consume two switch IDs.

C. Emulated switch IDs must be numbered from 1 to 99.

D. Each switch ID must be unique in the FabricPath topology.

E. Switch IDs must be configured manually.

Answer: B,D,E

Explanation:

To understand this feature, please refer to the link given below.

Reference: http://www.cisco.com/c/en/us/products/collateral/switches/nexus-5000-series-switches/guide_c07-690079.html#wp9000065

8. 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. Follow these guidelines when you upgrade or downgrade EPLDs:

•You can execute an upgrade from the active supervisor module only. All the modules, including the active supervisor module, can be updated individually.

•You can individually update each module whether it is online or offline as follows:

–If you upgrade EPLD images on an online module, only the EPLD images with version numbers that differ from the new EPLD images are upgraded.

–If you upgrade EPLD images on an offline module, all of the EPLD images are upgraded.

•On a system that has two supervisor modules, upgrade the EPLDs for the standby supervisor and then switch the active supervisor to standby mode to upgrade its EPLDs. On a system that has only one supervisor module, you can upgrade the active supervisor, but this will disrupt its operations during the upgrade.

•If you interrupt an upgrade, you must upgrade the module that is being upgraded again.

•The upgrade process disrupts traffic on the targeted module.

9. Which statement about Cisco FabricPath is true?

A. It is the best solution for interconnecting multiple data centers.

B. It optimizes STP throughout the Layer 2 network.

C. It is a simplified extension of Layer 3 networks across a single data center.

D. The Cisco FabricPath domain appears as a single STP bridge, where each edge port uses the same MAC address.

Answer: D

Explanation:

To have a loop-free topology for the CE/FabricPath hybrid network, the FabricPath network automatically displays as a single bridge to all connected CE devices. The STP domains do not cross into the FabricPath network. If multiple STP domains are defined, BPDUs and topology change notifications (TCNs) are localized to the domain. If a connected STP domain is multihomed to the FabricPath domain, a TCN must be able to reach to all devices in the STP domain through the FabricPath domain. As a result, the TCN is sent to the FabricPath domain through the IS-IS protocol data unit (PDU) by default.

Reference: http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/mkt_ops_guides/513_n1_1/n5k_ops_fabricpath.html

10. Refer to the exhibit.

```
OTV_EDGE1_SITE#1 show otv route
 OTV Unicast MAC Routing Table For Overlay1

VLAN MAC-Address          Metric Uptime  Last Updt   Owner
    Next-Hop(s)
!100 MACs from SITE 1 - local
110 0000.6e01.010a 1      2d16h          2d16h       lmac
    port-channel1


!100 MACs from SITE 2
110 0000.6e02.020a 42  2d16h          2d16h       isis_otv-default
    Overlay1-10.3.8.2


OTV_EDGE1_SITE#1 show otv route
 OTV Unicast MAC Routing Table For Overlay1

VLAN MAC-Address          Metric Uptime  Last Updt   Owner
    Next-Hop(s)
!100 MACs from SITE 1 - local
110 0000.6e01.010a 1      3d16h          3d16h       lmac
    port-channel1
110 0000.6e02.020a 1      0d01h          0d01h       lmac
    port-channel2

!100 MACs from SITE 2
```

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

11. What mode is required on a Cisco Nexus 7000 32-port 10-GB module port group to allow equal access to the 10-GB port controller?

A. dedicated

B. assigned

C. shared

D. community

Answer:: C

Explanation:

You can share 10 Gb of bandwidth among a group of ports (four ports) on a 32-port 10- Gigabit Ethernet module. To share the bandwidth, you must bring the dedicated port administratively down, specify the ports that are to share the bandwidth, change the rate mode to shared, and then bring the ports administratively up.

Reference: http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/interfaces/configuration/guide/if_cli/if_basic.html#70242

12. Which statement about scalability in Cisco OTV is true?

A. The control plane avoids flooding by exchanging MAC reachability.

B. IP-based functionality provides Layer 3 extension over any transport.

C. Any encapsulation overhead is avoided by using IS-IS.

D. Unknown unicasts are handled by the authoritative edge device.

Answer: A

Explanation:

Cisco calls the underlying concept of OTV traffic forwarding "MAC routing", since it behaves as if you are routing Ethernet frames over the DCI transport. OTV uses a control plane protocol to proactively propagate MAC address reachability before traffic is allowed to pass, which eliminates dependency on flooding mechanism to either learn MAC addresses or forward unknown unicasts.

Reference: http://www.computerworld.com/article/2515468/data-center/layer-2-data-center-interconnect-options.html

13. Refer to the exhibit.

```
Nexus# show glbp
Ethernet2/6 – Group 1
State is Up
1 state change(s), last state change(s)
00:02:53
Virtual IP address is 10.1.2.7
Hello time 3 sec, hold time 10 sec
Redirect time 600 sec, forwarded time-out
14400 sec
Preemption disabled
Active is unknown
Standby is unknown
Priority 100 (configured)
Weighting 100 (configured 100),
Thresholds: lower 1, upper 100
Load balancing: round-robin
Group members:
0015.1758.19AE (10.1.2.6) local
There are no forwarders
```

This multilayer Cisco Nexus switch had been the active virtual gateway for Group 1 before it became temporarily unavailable. What will happen to GLBP Group 1 when this device becomes available again?

A. The currently active router remains active.

B. It depends on the priority value that is configured active on the router.

C. The Cisco Nexus switch becomes the active virtual gateway after 600 seconds.

D. It depends on the weighting values that are configured active on the router.

Answer: A

Explanation:

GLBP prioritizes gateways to elect an active virtual gateway (AVG). If multiple gateways have the same priority, the gateway with the highest real IP address becomes the AVG. The AVG assigns a virtual MAC address to each member of the GLBP group. Each member is the active virtual forwarder (AVF) for its assigned virtual MAC address, forwarding packets sent to its assigned virtual MAC address.

The AVG also answers Address Resolution Protocol (ARP) requests for the virtual IP address. Load sharing is achieved when the AVG replies to the ARP requests with different virtual MAC addresses.

Note: Packets received on a routed port destined for the GLBP virtual IP address terminate on the local router, regardless of whether that router is the active GLBP router or a redundant GLBP router. This termination includes ping and Telnet traffic. Packets received on a Layer 2 (VLAN) interface destined for the GLBP virtual IP address terminate on the active router.

14. What must be enabled on the interface of a multicast-enabled device to support the Source Specific Multicast feature?

A. IGMP version 3

B. IGMP version 2

C. IGMP version 1

D. PIM

Answer: A

Explanation:

IGMP is the Internet Engineering Task Force (IETF) standards track protocol used for hosts to signal multicast group membership to routers. Version 3 of this protocol supports source filtering, which is required for SSM. To run SSM with IGMPv3, SSM must be supported in the Cisco IOS router, the host where the application is running, and the application itself. IGMP v3lite and URD are two Cisco-developed transition solutions that enable the immediate development and deployment of SSM services, without the need to wait for the availability of full IGMPv3 support in host operating systems and SSM receiver applications. IGMP v3lite is a solution for application developers that allows immediate development of SSM receiver applications switching to IGMPv3 as soon as it becomes available. URD is a solution for content providers and content aggregators that enables them to deploy receiver applications that are not yet SSM enabled (through support for IGMPv3). IGMPv3, IGMP v3lite, and URD interoperate with each other, so that both IGMP v3lite and URD can easily be used as transitional solutions toward full IGMPv3 support in hosts.

Reference: http://www.cisco.com/c/en/us/td/docs/ios/12_2/ip/configuration/guide/fipr_c/1cfssm.html

15. What is an Overlay Transport Virtualization extended VLAN?

A. the VLAN used to locate other AEDs

B. the VLAN used to access the overlay network by the join interface

C. the user VLAN that exists in multiple sites

D. the VLAN that must contain the overlay interface

Answer: C

Explanation: Functions of OTV

? Maintains a list of overlays

? Maintains a list of configured overlay parameters such as name, multicast address, encapsulation type, authentication, and OTV feature sets

? Maintains the state of the overlay interface

? Maintains the status of OTV VLAN membership from Ethernet infrastructure and the state of the authoritative edge device (AED) from IS-IS

? Maintains a database of overlay adjacencies as reported by IS-IS

? Maintains IP tunnel information and manages the encapsulation for data sent on the overlay network

? Manages delivery groups (DGs) for each overlay by snooping multicast traffic and monitoring traffic streams for active DGs

? Configures, starts, and stops the OTV IS-IS instance

? Interfaces with IP multicast to join provider multicast groups for each overlay

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 642-997 Exam with Our Prep Materials Via below:

http://www.certleader.com/642-997-dumps.html