Cisco.Premium.300-165.ny.VCEplus.195q

Exam Code: 300-165
Exam Name: Implementing Cisco Data Center Infrastructure (DCII)
Certification Provider: Cisco
Corresponding Certifications: CCNP Data Center
Website: www.vceplus.com
Free Exam: https://vceplus.com/ccnp-300-165/

Questions & Answers Exam Engine is rigorously checked before being put up for sale. We make sure there is nothing irrelevant in 300-165 exam products and you get latest questions. We strive to deliver the best 300-165 exam product for top grades in your first attempt.

VCE to PDF Converter : https://vceplus.com/vce-to-pdf/
Facebook: https://www.facebook.com/VCE.For.All.VN/
Twitter : https://twitter.com/VCE_Plus
Google+ : https://plus.google.com/+Vcepluscom
LinkedIn : https://www.linkedin.com/company/vceplus

QUESTION 1
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.

Correct Answer: D
Section: (none)
Explanation

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

QUESTION 2
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.

Correct Answer: A
Section: (none)
Explanation

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

QUESTION 3
Which two statements about Cisco Nexus 7000 line cards are true? (Choose two.)

A. M1, M2, and F1 cards are allowed in the same VDC.

B. M line cards are service-oriented and likely face the access layer and provide Layer 2 connectivity.

C. F line cards are performance-oriented and likely connect northbound to the core layer for Layer 3 connectivity.

D. M line cards support Layer 2, Layer 3, and Layer 4 with large forwarding tables and a rich feature set.

E. The F2 line card must reside in the admin VDC.

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Cisco is introducing a new line card called as F3 Module which has rich feature set and offers high performance 40G/100G port density to the Nexus 7000 product family. Cisco also introduced a new feature in NX-OS 6.2(2) where the F2e line card can be in the same VDC as M1 or M2

Line Card. The objective of this session is to cover detailed steps and methodology of migrating Nexus 7000 with VDC types prior to NX-OS 6.2 to the newer F3 or M/F2e VDC types. The session also covers the effect of VDC migration with commonly used Network features, firewall and load balancer services.
M-Series XL modules support larger forwarding tables. M-Series modules are frequently required at network core, peering, and aggregation points. When used with the F1-Series, the M-Series modules provide inter-VLAN services and form a pool of Layer 3 resources for the system. https://www.ciscolive2014.com/ connect/sessionDetail.ww?SESSION_ID=2244 http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Data_Center/VMDC/26/vmdctechwp.html

QUESTION 4
Which statement about the Layer 3 card on the Cisco Nexus 5500 Series Switch is true?

A. BGP support is not provided, but RIP, EIGRP, and OSPF support is provided.

B. Up to two 4-port cards are supported with up to 160 Gb/s of Layer 3 forwarding capability.

C. Up to 16 FEX connections are supported.

D. Port channels cannot be configured as Layer 3 interfaces.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

From the Cisco NX-OS 5.1(3)N1(1) release and later releases, each Cisco Nexus 5500 Series device can manage and support up to 24 FEXs without Layer 3. With Layer 3, the number of FEXs supported per Cisco Nexus 5500 Series device is 8. With Enhanced vPC and a dualhomed FEX topology each FEX is managed by both Cisco Nexus 5000 Series devices. As a result, one pair of Cisco Nexus 5500 Series devices can support up to 24 FEXs and 16 FEXs for Layer 2 and Layer 3. http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/mkt_ops_guides/513_n1_1/ n5k_enhanced_vpc.html

QUESTION 5
Which statement explains why a Cisco UCS 6200 Fabric Interconnect that is configured in end- host mode is beneficial to the unified fabric network?

A. There is support for multiple (power of 2) uplinks.

B. Upstream Layer 2 disjoint networks will remain separated.

C. The 6200 can connect directly via vPC to a Layer 3 aggregation device.

D. STP is not required on the uplink ports from the 6200.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In Cisco Unified Computing System environments, two Ethernet switching modes determine the way that the fabric interconnects behave as switching devices between the servers and the network. In end-host mode, the fabric interconnects appear to the upstream devices as end hosts with multiple links. In end-host mode, the switch does not run Spanning Tree Protocol and avoids loops by following a set of rules for traffic forwarding. In switch mode, the switch runs Spanning Tree Protocol to avoid loops, and broadcast and multicast packets are handled in the traditional way.
http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unifiedcomputing/whitepaper_c11-701962.html

QUESTION 6
Which option is a restriction of the unified ports on the Cisco UCS 6200 Series Fabric Interconnect when connecting to the unified fabric network?

A. Direct FC connections are not supported to Cisco MDS switches

B. The FCoE or Fibre Channel port allocations must be contiguous on the 6200.

C. 10-G Fibre Channel ports only use SFP+ interfaces.

D. vPC is not supported on the Ethernet ports.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
When you configure the links between the Cisco UCS 2200 Series FEX and a Cisco UCS 6200 series fabric interconnect in fabric port channel mode, the available

VIF namespace on the adapter varies depending on where the FEX uplinks are connected to the fabric interconnect ports.
Inside the 6248 fabric interconnect there are six sets of eight contiguous ports, with each set of ports managed by a single chip. When uplinks are connected such that all of the uplinks from an
Cisco 642-997 Exam
FEX are connected to a set of ports managed by a single chip, Cisco UCS Manager maximizes the number of VIFs used in service profiles deployed on the blades in the chassis. If uplink connections from an IOM are distributed across ports managed by separate chips, the VIF count is decreased.
http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/hw/6200-installguide/6200_HIG/6200_HIG_chapter_01.html

QUESTION 7
Which statement about the implementation of Cisco TrustSec on Cisco Nexus 7000 Series Switches is true?

A. While SGACL enforcement and SGT propagation are supported on the M and F modules, 802.1AE (MACsec) support is available only on the M module.

B. SGT Exchange Protocol is required to propagate the SGTs across F modules that lack hardware support for Cisco TrustSec.

C. AAA authentication and authorization is supported using TACACS or RADIUS to a Cisco Secure Access Control Server.

D. Both Cisco TrustSec and 802.1X can be configured on an F or M module interface.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The M-Series modules on the Nexus 7000 support 802.1AE MACSEC on all ports, including the new M2-series modules. The F2e modules will have this feature enabled in the future. It is important to note that because 802.1AE MACSEC is a link-level encryption, the two MACSECenabled endpoints, Nexus 7000 devices in our case, must be directly L2 adjacent.
This means we direct fiber connection or one facilitated with optical gear is required. MACSEC has integrity checks for the frames and intermediate devices, like another switch, even at L2, will cause the integrity checks to fail. In most cases, this means metro-Ethernet services or carrierprovided label switched services will not work for a MACSEC connection.
http://www.ciscopress.com/articles/article.asp?p=2065720

QUESTION 8
Which statement about implementation of Cisco TrustSec on Cisco Nexus 5546 or 5548 switches are true?

A. Cisco TrustSec support varies depending on Cisco Nexus 5500 Series Switch model.

B. The hardware is not able to support MACsec switch-port-level encryption based on IEEE 802.1AE.

C. The maximum number of RBACL TCAM user configurable entries is 128k.

D. The SGT Exchange Protocol must use the management (mgmt 0) interface.

Correct Answer: B
Section: (none)

Explanation

Explanation/Reference:
Explanation: https://scadahacker.com/library/Documents/Manuals/Cisco%20-%20TrustSec%20Solution%20O verview.pdf

QUESTION 9
Which two security features are only supported on the Cisco Nexus 7000 Series Switches? (Choose two.)

A. IP source guard
B. traffic storm control
C. CoPP
D. DHCP snooping
E. Dynamic ARP Inspection
F. NAC

Correct Answer: BF
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A traffic storm occurs when packets flood the LAN, creating excessive traffic and degrading network performance. You can use the traffic storm control feature to prevent disruptions on Layer 2 ports by a broadcast, multicast, or unicast traffic storm on physical interfaces. Traffic storm control (also called traffic suppression) allows you to monitor the levels of the incoming broadcast, multicast, and unicast traffic over a 10-millisecond interval. During this interval, the traffic level, which is a percentage of the total available bandwidth of the port, is compared with the traffic storm control level that you configured. When the ingress traffic reaches the traffic storm control level that is configured on the port, traffic storm control drops the traffic until the interval ends.
http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/dcnm/security/configuration/guide/b_Cisco_DCNM_Security_Configuration_Guide__Release_5x/ Cisco_DCNM_Security_Configuration_Guide__Release_5-x_chapter17.html
http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/dcnm/security/configuration/guide/b_Cisco_DCNM_Security_Configuration_Guide__Release_5-x/Cisco_DCNM_Security_Configuration_Guide__Release_5-x_chapter1.html

QUESTION 10
After enabling strong, reversible 128-bit Advanced Encryption Standard password type-6 encryption on a Cisco Nexus 7000, which command would convert existing plain or weakly encrypted passwords to type-6 encrypted passwords?

A. switch# key config-key ascii
B. switch(config)# feature password encryption aes
C. switch# encryption re-encrypt obfuscated
D. switch# encryption decrypt type6

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
This command converts existing plain or weakly encrypted passwords to type-6 encrypted
passwords. http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nxos/security/configuration/guide/b_Cisco_Nexus_7000_NX-OS_Security_Configuration_Guide__Release_5-x/b_Cisco_Nexus_7000_NXOS_Security_Configuration_Guide__Release_5-x_chapter_010101.html

QUESTION 11
By default it will take 10 seconds for authentication to fail due to an unresponsive RADIUS server before a Cisco Nexus series switch reverts to another RADIUS server or local authentication. What is one efficient way to improve the reaction time to a RADIUS server failure?

A. Decrease the global RADIUS retransmission count to 1.
B. Decrease the global RADIUS timeout interval to 5 seconds.
C. Configure the RADIUS retransmission count and timeout interval per server, versus globally.
D. Configure per server a test idle timer, along with a username and password.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You can monitor the availability of RADIUS servers. These parameters include the username and password to use for the server and an idle timer. The idle timer specifies the interval during which a RADIUS server receives no requests before the Nexus 5000 Series switch sends out a test packet. You can configure this option to test servers periodically. The test idle timer specifies the interval during which a RADIUS server receives no requests before the Nexus 5000 Series switch sends out a test packet. The default idle timer value is 0 minutes. When the idle time interval is 0 minutes, the Nexus 5000 Series switch does not perform periodic RADIUS server monitoring. http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli_r el_4_0_1a/CLIConfigurationGuide/sec_radius.html

QUESTION 12
Which statement about RADIUS configuration distribution using Cisco Fabric Services on a Cisco Nexus 7000 Series Switch is true?

A. Cisco Fabric Services does not distribute the RADIUS server group configuration or server and global keys.
B. Enabling Cisco Fabric Services causes the existing RADIUS configuration on your Cisco NX- OS device to be immediately distributed.
C. When the RADIUS configuration is being simultaneously changed on more than one device in a Cisco Fabric Services region, the most recent changes will take precedence.

D. Only the Cisco NX-OS device with the lowest IP address in the Cisco Fabric Services region can lock the RADIUS configuration.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
CFS does not distribute the RADIUS server group configuration or server and global keys. The keys are unique to the Cisco NX-OS device and are not shared with other Cisco NX-OS devices.
http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/6_x/nxos/security/configuration/guide/b_Cisco_Nexus_7000_NX-OS_Security_Configuration_Guide__Release_6-x/b_Cisco_Nexus_7000_NXOS_Security_Configuration_Guide__Release_6-x_chapter_0101.html

QUESTION 13
When a local RBAC user account has the same name as a remote user account on an AAA server, what happens when a user with that name logs into a Cisco Nexus switch?

A. The user roles from the remote AAA user account are applied, not the configured local user roles.

B. All the roles are merged (logical OR).

C. The user roles from the local user account are applied, not the remote AAA user roles.

D. Only the roles that are defined on both accounts are merged (logical AND).

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If you have a user account configured on the local Cisco NX-OS device that has the same name as a remote user account on an AAA server, the Cisco NX-OS software applies the user roles for the local user account to the remote user, not the user roles configured on the AAA server. http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/4_1/nxos/security/configuration/guide/sec_nx-os-cfg/sec_rbac.html

QUESTION 14
Which statement is true if password-strength checking is enabled?

A. Short, easy-to-decipher passwords will be rejected.

B. The strength of existing passwords will be checked.

C. Special characters, such as the dollar sign ($) or the percent sign (%), will not be allowed.

D. Passwords become case-sensitive.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If a password is trivial (such as a short, easy-to-decipher password), the cisco NX_OS software will reject your password configuration if password-strength checking is enabled. Be sure to configure a strong password. Passwords are case sensitive.
http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7x/security/configuration/guide/
b_Cisco_Nexus_9000_Series_NXOS_Security_Configuration_Guide_7x/
b_Cisco_Nexus_9000_Series_NXOS_Security_Configuration_Guide_7x_chapter_01000.pdf

QUESTION 15
Which statement about RBAC user roles on a Cisco Nexus switch is true?

A. If you belong to multiple roles, you can execute only the commands that are permitted by both roles (logical AND).
B. Access to a command takes priority over being denied access to a command.
C. The predefined roles can only be changed by the network administrator (superuser).
D. The default SAN administrator role restricts configuration to Fibre Channel interfaces.
E. On a Cisco Nexus 7000 Series Switch, roles are shared between VDCs.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If you belong to multiple roles, you can execute a combination of all the commands permitted by these roles. Access to a command takes priority over being denied access to a command. For example, suppose a user has RoleA, which denied access to the configuration commands. However, the users also have RoleB, which has access to the configuration commands. In this case, the users have access to the configuration commands. http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/C LIConfigurationGuide/sec_rbac.html

QUESTION 16
Which two statements about SAN zoning on Cisco Nexus switches are true? (Choose two.)

A. Unlike configured zones, default zone information is not distributed to the other switches in the fabric. B. Traffic can either be permitted or denied among members of the default zone. This information is not distributed to all switches. It must be configured in each switch.
B. The settings for default zone configurations cannot be changed.
C. To activate a zone set, you must copy the running configuration to the startup configuration after the zone set is configured.

D. Soft zoning restrictions will not prevent a source device from accessing a device outside its zone, if the source knows the Fibre Channel ID of the destination.

E. Hard zoning is enforced by the hardware on each FLOGI sent by an N Port.

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Each member of a fabric (in effect a device attached to an Nx port) can belong to any zone. If a member is not part of any active zone, it is considered to be part of the default zone. Therefore, if no zone set is active in the fabric, all devices are considered to be in the default zone. Even though a member can belong to multiple zones, a member that is part of the default zone cannot be part of any other zone. The switch determines whether a port is a member of the default zone when the attached port comes up. Unlike configured zones, default zone information is not distributed to the other switches in the fabric Traffic can either be permitted or denied among members of the default zone. This information is not distributed to all switches; it must be configured in each switch. http://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/5_2/configuration/guides/f abric/DCNM-SAN/fm_fabric/zone.html

QUESTION 17
Which two statements about SAN zoning on Cisco Nexus switches are true? (Choose two.)

A. Zoning is enforced by examining the destination ID field.

B. Devices can only belong to one zone.

C. Only one zone set can be activated at any time.

D. A zone can only be a member one zone set.

E. Zoning must be administered from the primary SAN switch in the fabric.

F. Zone configuration changes are nondisruptive.

Correct Answer: CF
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A zone set can be activated or deactivated as a single entity across all switches in the fabric. Only one zone set can be activated at any time. If zoning is not activated, all devices are members of the default zone. If zoning is activated, any device that is not in an active zone (a zone that is part of an active zone set) is a member of the default zone. Zoning can be administered from any switch in the fabric. When you activate a zone (from any switch), all switches in the fabric receive the active zone set. Additionally, full zone sets are distributed to all switches in the fabric, if this feature is enabled in the source switch.
http://www.cisco.com/en/US/docs/storage/san_switches/mds9000/sw/sanos/quick/guide/qcg_zones.html

QUESTION 18
The Connectivity Management Processor monitors the active supervisor module on a Cisco Nexus 7000 switch and will reboot the device in the event of a lights-

out management issue. However, which option includes features that provide similar benefits in the absence of the Connectivity Management Processor?

A. high-availability functionality from features such as vPC and NSF
B. traditional system connectivity models like SNMP, GUI, or SSH
C. Cisco FabricPath
D. VDC failover

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation: vPC uses the vPC peer-keepalive link to run hello messages that are used to detect a dual-active scenario. A Gigabit Ethernet port can be used to carry the peer-keepalive messages. A dedicated VRF is recommended to isolate these control messages from common data packets. When an out-of-band network infrastructure is present, the management interfaces of the Cisco Nexus 7000 supervisor could be also used to carry keep-alive connectivity using the dedicated management VRF. When the vPC peer-link is no longer detected, a dual-active situation occurs, and the system disables all vPC port channel member on the "secondary" vPC peer (lower vPC role priority value). Also SVI interfaces associated to a vPC VLAN are suspended on the secondary switch. As a result, in this condition only the "primary" vPC peer actively forwards traffic on the vPC VLANs. Multiple peer-keepalive links can be used to increase resiliency of the dual-active detection mechanism.
Both the Cisco Catalyst 6500 and the Cisco Nexus 7000 offer a variety of high-availability features. Some of the primary features to highlight are In Service Software Upgrade (ISSU), Stateful Switchover (SSO), and Nonstop Forwarding (NSF). The operation and the behavior of these features are unique to the respective platform and can be independently executed without affecting the interoperability between the two platforms.
http://www.cisco.com/c/en/us/products/collateral/switches/catalyst-6500-seriesswitches/white_paper_c11_589890.html

QUESTION 19
Which Cisco Nexus feature is best managed with DCNM-SAN?

A. VSS
B. domain parameters
C. virtual switches
D. AAA

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Fibre Channel domain (fcdomain) feature performs principal switch selection, domain
ID distribution, FC ID allocation, and fabric reconfiguration functions as described in the FC-SW-2 standards. The domains are configured on a per VSAN basis. If

you do not configure a domain ID, the local switch uses a random ID.
This section describes each fcdomain phase:
Principal switch selection - This phase guarantees the selection of a unique principal switch across the fabric.
Domain ID distribution - This phase guarantees each switch in the fabric obtains a unique domain ID.
FC ID allocation - This phase guarantees a unique FC ID assignment to each device attached to the corresponding switch in the fabric.
Fabric reconfiguration - This phase guarantees a resynchronization of all switches in the fabric to ensure they simultaneously restart a new principal switch selection phase. http://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/5_2/configuration/guides/ sysmgnt/DCNM-SAN/sysmgmt_dcnm/ sysmgmt_overview.html#wp1051962

QUESTION 20
Which of the following Cisco Nexus features is best managed with DCNM-LAN?

A. VSS
B. Domain parameters
C. Virtual switches
D. AAA

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

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.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
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

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.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
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/nxos/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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
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/nxos/unicast/configuration/guide/l3_cli_nxos/l3_glbp.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

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
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

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:
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-seriesswitches/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.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
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_op s.html

QUESTION 27
Which policy-map action performs congestion avoidance?

A. priority
B. bandwidth
C. queue-limit
D. random-detect

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
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/qcfconav.html

QUESTION 28
Refer to the exhibit. Which statement based on these two outputs that were collected 24 hours apart is true?

```
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

```

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.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

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.

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:
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-catmulticast-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.

Correct Answer: CD
Section: (none)
Explanation

Explanation/Reference:
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.

QUESTION 31
Which three VDC resources can be constrained with a resource template? (Choose three.)

A. ACLs
B. NAT entries
C. IPv4 routes
D. IPv6 routes
E. SPAN sessions
F. RBAC users

Correct Answer: CDE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
VDC resource templates set the minimum and maximum limits for shared physical device resources when you create the VDC. The Cisco NX-OS software reserves the minimum limit for the resource to the VDC. Any resources allocated to the VDC beyond the minimum are based on the maximum limit and availability on the device.
You can explicitly specify a VDC resource template, or you can use the default VDC template provided by the Cisco NX-OS software. VDC templates set limits on the following resources:
IPv4 multicast route memory
IPv6 multicast route memory
IPv4 unicast route memory
IPv6 unicast route memory
Port channels
Switch Port Analyzer (SPAN) sessions
VLANs
Virtual routing and forwarding instances (VRFs)

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nx-os/virtual_device_context/configuration/guide/b-7k-Cisco-Nexus-7000-Series-NX-OS-VirtualDevice-Context-Configuration-Guide/vdc-res-template.html

QUESTION 32
Which command sequence correctly enables Adapter FEX on Nexus 5000 Series Switches?

A. switch(config)# install feature-set virtualization switch(config)# feature-set virtualization
B. switch(config)# install feature-set adapter-fex switch(config)# feature-set adapter-fex
C. switch(config)# install feature-set adapter-fex switch(config)# feature-set virtualization
D. switch(config)# install feature-set virtualization switch(config)# feature-set adapter-fex

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
install feature-set virtualization : installs the cisco virtual machine feature set on the switch. feature-set virtualization : enables the cisco virtual machine feature on the switch.
http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/adapterfex/513_n1_1/
b_Configuring_Cisco_Nexus_5000_Series_AdapterFEX_rel_5_1_3_N1/b_Configuring_Cisco_Nexus_5000_Series_Adapter-FEX_rel_5_1_3_N1_chapter_010.pdf

QUESTION 33
Which three Cisco UCS C-Series CNAs support Adapter FEX? (Choose three.)

A. Qlogic QLE8152
B. Broadcom BCM57712
C. Cisco UCS P81E
D. Cisco UCS VIC 1220
E. Emulex OCe10102-FX-C
F. Intel X520

Correct Answer: BCD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/c-series_integration/ucsm2-1/b_UCSM2-1_C-Integration/b_UCSM2-1_C-

Integration_chapter_011.html#reference_D644111FC68046F0BEA49756A0834664

QUESTION 34
Which two Cisco Nexus platforms support Adapter FEX? (Choose two.)

A. Cisco Nexus 7000 Series Switches B. Cisco Nexus 5000 Series Switches

B. Cisco Nexus 5500 Series Switches

C. Cisco Nexus 4000 Series Switches

D. Cisco Nexus 2000 Series Fabric Extenders

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
At the access layer, the Adapter-FEX requires a FEX-enabled adapter on a server that connects to a parent device that supports virtualization of interfaces. The Adapter-FEX is supported on the following platforms:
• The Cisco Unified Computing System (UCS) platform supports Adapter-FEX between UCS servers and the UCS Fabric Interconnect.
• The Adapter-FEX is supported on the Cisco Nexus 5500 Series platform and on the Cisco Nexus 2200 Fabric Extender that is connected to a Cisco Nexus 5500 Series parent device. This implementation works on a variety of FEX-capable adapters, including the Cisco UCS P81E virtual interface card (VIC) adapter for the UCS C-Series platform and third party adapters such as the Broadcom BCM57712 Convergence Network Interface Card, that implement the virtual network tag (VNTag) technology.
http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/operations/adapter_fex/ 513_n1_1/ops_adapter_fex/ops_using_adapter_fex.html

QUESTION 35
Which three items must be configured in the port profile client in Cisco UCS Manager? (Choose three.)

A. port profile

B. DVS

C. data center

D. folder

E. vCenter IP address

F. VM port group

Correct Answer: BCD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
After associating an ESX host to a DVS, you can migrate existing VMs from the vSwitch to the
DVS, and you can create VMs to use the DVS instead of the vSwitch. With the hardware-based
VN-Link implementation, when a VM uses the DVS, all VM traffic passes through the DVS and ASIC-based switching is performed by the fabric interconnect. In Cisco UCS Manager, DVSes are organized in the following hierarchy:
vCenter Folder (optional)
Datacenter
Folder (required)
DVS
At the top of the hierarchy is the vCenter, which represents a VMware vCenter instance. Each vCenter contains one or more datacenters, and optionally vCenter folders with which you can organize the datacenters. Each datacenter contains one or more required datacenter folders.
Datacenter folders contain the DVSes.
http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/gui/config/guide/1-3-1/b_UCSM_GUI_Configuration_Guide_1_3_1/
UCSM_GUI_Configuration_Guide_1_3_1_chapter 2 8.html

QUESTION 36
In the dynamic vNIC creation wizard, why are choices for Protection important?

A. They allow reserve vNICs to be allocated out of the spares pool.
B. They enable hardware-based failover.
C. They select the primary fabric association for dynamic vNICs.
D. They allow dynamic vNICs to be reserved for fabric failover.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Number of Dynamic vNICs - This is the number of vNICs that will be available for dynamic assignment to VMs. Remember that the VIC has a limit to the number of vNICs that it can support and this is based on the number of uplinks between the IOM and the FI. At least this is the case with the 2104 IOM and the M81KR VIC, which supports ((# IOM Links * 15) ?2)). Also remember that your ESXi server will already have a number of vNICs used for other traffic such as Mgmt, vMotion, storage, etc, and that these count against the limit.
Adapter Policy - This determines the vNIC adapter config (HW queue config, TCP offload, etc) and you must select VMWarePassThru to support VM-FEX in High Performance mode.
Protection - This determines the initial placement of the vNICs, either all of them are placed on fabric A or Fabric B or they are alternated between the two fabrics if you just select the "Protected" option. Failover is always enabled on these vNICs and there is no way to disable the protection.
http://infrastructureadventures.com/2011/10/09/deploying-cisco-ucs-vm-fex-forvsphere-%E2%80%93-part-2-ucsm-config-and-vmware-integration/

QUESTION 37

How is a dynamic vNIC allocated?

A. Dynamic vNICs are assigned to VMs in vCenter.
B. Dynamic vNICs can only be bound to the service profile through an updating template.
C. Dynamic vNICs are bound directly to a service profile.
D. Dynamic vNICs are assigned by binding a port profile to the service profile.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The dynamic vNIC connection policy determines how the connectivity between VMs and dynamic vNICs is configured. This policy is required for Cisco UCS domains that include servers with VIC adapters on which you have installed VMs and configured dynamic vNICs. Each dynamic vNIC connection policy includes an Ethernet adapter policy and designates the number of vNICs that can be configured for any server associated with a service profile that includes the policy. For VM-FEX that has all ports on a blade in standard mode, you need to use the VMware adapter policy.
For VM-FEX that has at least one port on a blade in high-performance mode, use the VMwarePassThrough adapter policy or create a custom policy. If you need to create a custom policy, the resources provisioned need to equal the resource requirements of the guest OS that needs the most resources and for which you will be using high-performance mode.
http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/vm_fex/vmware/gui/config_guide/b_GUI_VMware_VM-FEX_UCSM_Configuration_Guide/
b_GUI_VMware_VMFEX_UCSM_Configuration_Guide_chapter_010.html

QUESTION 38
Refer to the command below. When configuring an SVS connection on the Cisco Nexus 5000 Series Switch, which device is being referenced as the remote IP address?
 nexus5500-2(config-svs-conn)# remote ip address 10.10.1.15 port 80 vrf management

A. ESX or ESXi host
B. vCenter
C. vPC peer switch
D. Cisco IMC management

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
This command specifies the hostname or IP address for the vCenter Server. Optionally, specifies the port number and VRF.

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5500/sw/layer2/6x/b_5500_Layer 2_Config_6x/
b_5500_Layer2_Config_602N12_chapter_010000.html

QUESTION 39
When connecting Cisco Nexus 5000 Series Switches to the VMware vCenter Server, which item must be configured before installing the extension keys?

A. configure vPC
B. configure DirectPath I/O support in vCenter
C. configure PTS on the VSM
D. configure dynamic vNICs

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 40
Which feature enables NIV?

A. EHV
B. vPC
C. Cisco FabricPath
D. Cisco OTV
E. VN-Tag

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
EHV is the feature that enables NIV.

QUESTION 41
Which three selections represent implementations of Cisco VN-Link technology? (Choose three.)

A. Cisco Nexus 1000V

B. Cisco Nexus 2000 FEX

C. Cisco VM-FEX

D. VMware PTS

E. vMotion

Correct Answer: ACD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The VM is powered on and resides on the ESX Host 1 with all the information stored on the shared storage.
The VM was connected to the PODy (where y is the number of your POD) PTS VDS by associating it to port group VLAN61 that was created on the Cisco Nexus 5548 device. The VM has been connected to the vPC system automatically using a VN-Link in the hardware in PTS mode or in VM-FEX mode.
The VEM bits are used in PTS mode to connect the VM VNIC to the VMNIC interface.
In this case, the VMNIC interface is not a real VMNIC but a dynamic VNIC that is presented as an interface to the ESX OS. The dynamic VNIC is enabled when the Cisco UCS VIC creates and configures the VNIC parameters inherited from port group VLAN61.
http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/mkt_ops_guides/513_n1_1/ n5k_ops_vmfex.html

QUESTION 42
Which two items are required components of VN-Link in software? (Choose two.)

A. VDC

B. VEM

C. vPC

D. VSM

E. VRRP

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Cisco Nexus 1000V Series consists of two main types of components that can virtually emulate a 66-slot modular Ethernet switch with redundant supervisor functions:
•Virtual Ethernet module (VEM)-data plane: This lightweight software component runs inside the hypervisor. It enables advanced networking and security features, performs switching between directly attached virtual machines, provides uplink capabilities to the rest of the network, and effectively replaces the vSwitch. Each hypervisor is embedded with one VEM.
•Virtual supervisor module (VSM)-control plane: This standalone, external, physical or virtual appliance is responsible for the configuration, management,

monitoring, and diagnostics of the overall Cisco Nexus 1000V Series system (that is, the combination of the VSM itself and all the VEMs it controls) as well as the integration with VMware vCenter. A single VSM can manage up to 64 VEMs. VSMs can be deployed in an active-standby model, helping ensure high availability. http://www.cisco.com/c/en/us/solutions/collateral/switches/nexus-1000v-switch-vmwarevsphere/white_paper_c11-525307.html

QUESTION 43
Which two items are features that are available in VN-Link in software? (Choose two.)

A. VM snapshot
B. NetFlow
C. ERSPAN
D. high availability
E. resource reservations

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:
Explanation:
NetFlow is a feature that was introduced on Cisco routers that provides the ability to collect IP network traffic as it enters or exits an interface. By analyzing the data provided by NetFlow, a network administrator can determine things such as the source and destination of traffic, class of service, and the causes of congestion. A typical flow monitoring setup (using NetFlow) consists of three main components:
▪ Flow exporter: aggregates packets into flows and exports flow records towards one or more flow collectors.
▪ Flow collector: responsible for reception, storage and pre-processing of flow data received from a flow exporter.
▪ Analysis application: analyzes received flow data in the context of intrusion detection or traffic profiling,

This module describes how to configure Encapsulated Remote Switched Port Analyzer (ERSPAN). The Cisco ERSPAN feature allows you to monitor traffic on one or more ports or VLANs and send the monitored traffic to one or more destination ports. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/lanswitch/configuration/xe-3s/lanswitch-xe-3sbook/lnsw-conf-erspan.html

QUESTION 44
Which statement about enhanced zoning on Cisco Multilayer Director Switches are true?

A. It allows partial zone set changes to be distributed without having to activate a zone set.
B. Enhanced zoning is compatible with IVR.
C. Zone changes can scheduled with a CRON job.
D. More than one zone set can be active with enhanced zoning.

Correct Answer: A
Section: (none)

Explanation

Explanation/Reference:
Explanation:
Enhanced zoning implements changes to the zoning database and distributes it without reactivation. Distribution of zone sets without activation avoids hardware changes for hard zoning in the switches.
http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5500/sw/san_switching/6x/b_550 0_SAN_Switching_Config_6x/
b_5500_SAN_Switching_Config_602N12_chapter_01001.html#con _1871274

QUESTION 45
Which command enables NPIV on Cisco Nexus 5000 Series Switches and Cisco MDS switches?

A. switch(config)# npiv enable

B. switch(config)# npivon

C. switch(config)# feature npiv

D. switch(config)# npiv proxy

E. switch(config)# np proxy-enable

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nxos/san_switching/configuration/guide/
b_Cisco_Nexus_7000_NXOS_SAN_Switching_Configuration_Guide/Cisco_Nexus_7000_NXOS_SAN_Switching_Configuration_Guide_chapter2.html

QUESTION 46
Between which two types of ports does FIP establish Fibre Channel virtual links? (Choose two.)

A. VE Ports and VE Ports

B. N Ports and F Ports

C. VN Ports and VF Ports

D. VP Ports and VE Ports

E. VE Ports and VF Ports

F. E Ports and E Ports

Correct Answer: AC
Section: (none)

Explanation

Explanation/Reference:
Explanation:
FIP aims to establish virtual FC links between VN_Ports and VF_Ports (ENode to FCF), as well as between pairs of VE_Ports (FCF to FCF), since these are the only legal combinations supported by native Fibre Channel fabrics. Standards-compliant implementations are not required to support both forms of virtual FC links, and Cisco has decided to focus initially on implementing FIP only between ENodes and FCFs. FCF-to-FCF connectivity is considered a strategic direction for end-to-end FCoE deployments, but the short-term urgency is for FCoE adoption between CNAs and the Fibre Channel fabric perimeter, where unified fabric can offer the greatest capital expenditure (CapEx) savings today. http://www.cisco.com/c/en/us/products/collateral/switches/nexus-7000-seriesswitches/white_paper_c11-560403.html

QUESTION 47
Which FCoE component is responsible for the encapsulation and de-encapsulation of Fibre Channel frames in Ethernet?

A. distributed FCF
B. FCoE node
C. FCoE logical endpoint
D. Fibre Channel forwarder
E. FCoE forwarder

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The FCoE Logical Endpoint (FCoE_LEP) is responsible for the encapsulation and deencapsulation functions of the FCoE traffic. FCoE_LEP has the standard Fibre Channel layers, starting with FC-2 and continuing up the Fibre Channel Protocol stack.
https://www.safaribooksonline.com/library/view/ccna-datacenter/9780133860429/ch11lev3sec5.html

QUESTION 48
Which item represents the process that allows FCoE multihop using T11 standard FC-BB-5?

A. distributed FCF
B. FIP proxy
C. N Port proxy
D. FIP snooping

Correct Answer: D

Section: (none)
Explanation

Explanation/Reference:
Explanation:
FIP snooping is used in multi-hop FCoE environments. FIP snooping is a frame inspection method that can be used by FIP snooping capable DCB devices to monitor FIP frames and apply policies based on the information in those frames. This allows for:
Enhanced FCoE security (Prevents FCoE MAC spoofing.)
Creates FC point-to-point links within the Ethernet LAN
Allows auto-configuration of ACLs based on name server information read in the FIP frames http://www.definethecloud.net/fcoe-initialization-protocol-fip-deep-dive/

QUESTION 49
How does an FCoE end node acquire its FCoE MAC address?

A. server-provided MAC address

B. Fibre Channel name server

C. fabric-provided MAC address

D. FIP proxy

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The VN_Port is assigned a fabric-provided Mac address (FPMA) that is built by concatenating a 24-bit FCoE MAC address prefix (FC-MAP), ranging from 0x0E-FC-00 to 0x0E-FC-FF, to the 24bit FCID. Being able to build a unique MAC address for the VN_Port directly from its FCID saves the switch from having to maintain a table that associates FCID and MAC addresses.
http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Data_Center/UF_FCoE_final.html

QUESTION 50
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

Correct Answer: C
Section: (none)

Explanation

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

QUESTION 51
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

Correct Answer: B
Section: (none)
Explanation

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

QUESTION 52
Which protocol is responsible for the discovery of FCoE capabilities on a remote switch?

A. DCE

B. DCBx

C. CDP

D. LLDP

Correct Answer: B
Section: (none)

Explanation

Explanation/Reference:
Explanation:
Data Center Bridging Capabilities Exchange Protocol (DCBX): a discovery and capability exchange protocol that is used for conveying capabilities and configuration of the above features between neighbors to ensure consistent configuration across the network. This protocol leverages functionality provided by IEEE 802.1AB (LLDP). It is actually included in the 802.1az standard.
http://en.wikipedia.org/wiki/Data_center_bridging

QUESTION 53
Which two items are services that are provided by Cisco Fabric Services? (Choose two.)

A. device alias distribution

B. VLAN database distribution

C. Kerberos proxy distribution

D. RSA key pair distribution

E. DPVM configuration distribution

Correct Answer: AE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The device alias application uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management and distribution. Device aliases use the coordinated distribution mode and the fabric-wide distribution scope.
DPVM can use CFS to distribute the database to all switches in the fabric. This allows devices to move anywhere and keep the same VSAN membership. You should enable CFS distribution on all switches in the fabric. Using the CFS infrastructure, each DPVM server learns the DPVM database from each of its neighboring switches during the ISL bring-up process. If you change the database locally, the DPVM server notifies its neighboring switches, and that database is updated by all switches in the fabric. http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/C LIConfigurationGuide/
ddas.html and http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nxos/san_switching/configuration/guide/
b_Cisco_Nexus_7000_NXOS_SAN_Switching_Configuration_Guide/Cisco_Nexus_7000_NX-
OS_SAN_Switching_Configuration_Guide_chapter4.html#concept_2B83E16506C845B39BDF96F9CAFFAEC3

QUESTION 54
On a Cisco Nexus 7000 Series router, which statement about HSRP and VRRP is true?

A. When VDCs are in use, only VRRP is supported.

B. HSRP and VRRP both use the same multicast IP address with different port numbers.

C. HSRP has shorter default hold and hello times.

D. The VRRP group IP address can be the same as the router-specific IP address.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
VRRP allows for transparent failover at the first-hop IP router by configuring a group of routers to share a virtual IP address. VRRP selects a master router in that group to handle all packets for the virtual IP address. The remaining routers are in standby and take over if the master router fails.
http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nxos/unicast/configuration/guide/l3_cli_nxos/l3_vrrp.html

QUESTION 55
Refer to the exhibit. 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?

```
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
```

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.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
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.

QUESTION 56
Which function does the graceful restart feature allow a Cisco Nexus 7000 Series router to perform?

A. Perform a rapid route convergence.

B. Initialize a standby supervisor transparently when one is present.

C. Remain in the data forwarding path through a process restart.

D. Maintain a management connection throughout a router restart.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Graceful Restart and Non Stop Routing both allow for the forwarding of data packets to continue along known routes while the routing protocol information is being restored (in the case of Graceful Restart) or refreshed (in the case of Non Stop Routing) following a processor switchover. When Graceful Restart is used, peer networking devices are informed, via protocol extensions prior to the event, of the SSO capable routers ability to perform graceful restart. The peer device must have the ability to understand this messaging. When a switchover occurs, the peer will continue to forward to the switching over router as instructed by the GR process for each particular protocol, even though in most cases the peering relationship needs to be rebuilt. Essentially, the peer router will give the switching over

router a "grace" period to re-establish the neighbor relationship, whilecontinuing to forward to the routes from that peer.
http://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/highavailability/solution_overview_c22-487228.html

QUESTION 57
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

Correct Answer: D
Section: (none)
Explanation

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

QUESTION 58
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

Correct Answer: A
Section: (none)
Explanation

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

QUESTION 59
Which two statements about implementing Cisco NPV and NPIV on a Cisco Nexus 5000 Series switch are true? (Choose two.)

A. STP must run inside the FP network.
B. All VLANs must be in the same mode, CE, or FP.
C. FP port can join the private and nonprivate VLANs.
D. Only F and M series modules can run FabricPath.
E. These require an enhanced Layer 2 license to run.

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
With the Nexus 5x00 switch, FCoE functionality is a licensed feature. After the license is installed, FCoE configuration can be completed. http://www.ciscopress.com/articles/article.asp?p=2030048&seqNum=4

QUESTION 60
What is the Overlay Transport Virtualization site VLAN used for?

A. to allow the join interfaces at different sites to communicate
B. to detect devices at the site that are not capable of OTV
C. to allow multiple site AEDs to communicate with each other
D. to detect other OTV edge devices in the site

Correct Answer: D
Section: (none)

Explanation

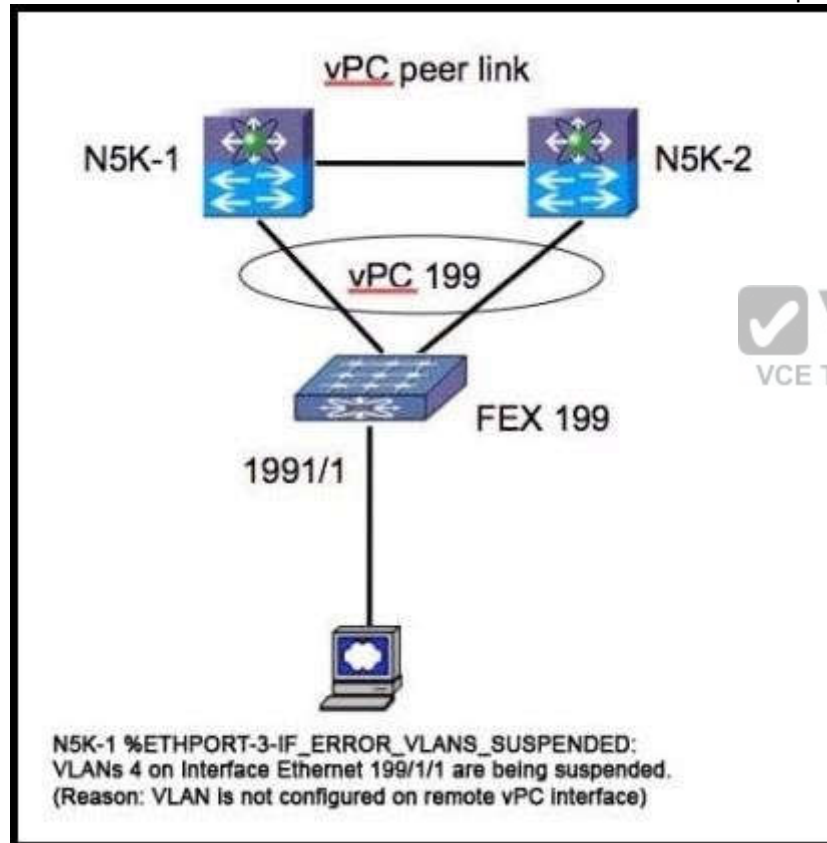**Explanation/Reference:**
Explanation:
The edge device performs OTV functions: it receives the Layer 2 traffic for all VLANs that need to be extended to remote locations and dynamically encapsulates the Ethernet frames into IP packets that are then sent across the transport infrastructure. It is expected that at least two OTV edge devices are deployed at each data center site to improve the resiliency. http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Data_Center/DCI/whitepaper/DCI3_OT V_Intro/DCI_1.html

QUESTION 61
Refer to the exhibit. Which corrective action is taken to resolve the problem?

A. Trunk four VLANs on interface ethernet 199/1/1.
B. Use the shut and no shut interface ethernet 199/1/1so that the VLANs come up.

C.   Place interface ethernet 199/1/1 in VLAN 4 in the N5K-2 configuration.
D.   Prune all but four VLANs from vPC 199.
E.   Add VLAN 4 to vPC 199.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Place interface ethernet 199/1/1 in VLAN 4 in the N5K-2 configuration.

QUESTION 62
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

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
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

QUESTION 63
Refer to the exhibit. What is the consequence of configuring peer-gateway on the two vPC peers N7K-1 and N7K-2?

```
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)#
```

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.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
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 noncomplaint with some basic Ethernet RFC standards. Packets reaching a vPC device for the nonlocal 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.
http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/4_2/nxos/interfaces/configuration/guide/if_nxos/if_vPC.html

QUESTION 64

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

```
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 No   Yes
  2    0022.5579.b1c2 Primary Confirmed No   Yes
  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
```

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.

Correct Answer: BDE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
To understand this feature, please refer to the link given below. http://www.cisco.com/c/en/us/products/collateral/switches/nexus-5000-seriesswitches/guide_c07-690079.html#wp9000065

QUESTION 65
Which statement about core-edge SAN topology is true?

A. Converged FCoE links connect the core and edge MDS switches.

B. The SAN core connects to the network aggregation layer.

C. Separate links with the same I/O are used for SAN and LAN traffic.

D. Storage devices are accessed via FCoE over the LAN network.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Aggregation layer of the data center provides connectivity for the Access layer switches in the server farm, an aggregates them into a smaller number of interfaces to be connected into the Core layer. In most data center environments, the Aggregation layer is the transition point between the purely Layer 3 routed Core layer, and the Layer 2-switched Access layer. 802.1Q trunks extend the server farm VLANs between Access and Aggregation layers. The Aggregation layer also provides a common connection point to insert services into the data flows between clients and servers, or between tiers of servers in a multi-tier application.

QUESTION 66
What configuration is required when implementing FCoE?

A. disable LAN traffic on the interface

B. configure PortFast on the access port

C. permit all VLANs on the interface

D. permit all VSANs on the interface

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
DCBX allows the switch to send a LAN Logical Link Status (LLS) message to a directlyconnected CNA. Enter the shutdown lan command to send an LLS-Down message to the CNA. This command causes all VLANs on the interface that are not enabled for FCoE to be brought down. If a VLAN on the interface is enabled for FCoE, it continues to carry SAN traffic without any interruption.
http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/fcoe/b_Cisco_Nexus_5
000_Series_NX-
OS_Fibre_Channel_over_Ethernet_Configuration_Guide_/
Cisco_Nexus_5000_Series_NXOS_Fibre_Channel_over_Ethernet_Configuration_Guide__chapter3.html

QUESTION 67
Which topology is not supported when using vPC?

A. a single-homed server to a single FEX that is connected to two Cisco Nexus 5500 Series Switches

B. a dual-homed server to two FEXs, each connected to two Cisco Nexus 5500 Series Switches

C. a dual-homed server to two FEXs that are connected to one Cisco Nexus 5500 Series Switch

D. a dual-homed server to a single FEX that is connected to two Cisco Nexus 5500 Series Switches

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The figure shows unsupported topology where a vPC is between hosts and two FEXs that are connected to one Cisco Nexus 5500 Series device. This topology does not provide a good high availability solution because the server loses the connectivity to the network when the Cisco Nexus 5000 Series device fails.
If you need to connect a multi-homing server to a pair of FEXs when there is only one Cisco Nexus 5000 Series device, you have the option to run active or standby NIC teaming from the server.
http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/mkt_ops_guides/513_n1_1/ n5k_enhanced_vpc.html

QUESTION 68
Which protocol is the foundation for unified fabric as implemented in Cisco NX-OS?

A. Fibre Channel

B. Data Center Bridging

C. Fibre Channel over Ethernet

D. N proxy virtualization

E. N Port identifier virtualization

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Fibre Channel over Ethernet (FCoE) is one of the major components of a Unified Fabric. FCoE is a new technology developed by Cisco that is standardized in the Fibre Channel Backbone 5 (FCBB-5) working group of Technical Committee T11 of the International Committee for Information
Technology Standards (INCITS). Most large data centers have huge installed bases of Fibre Channel and want a technology that maintains the Fibre Channel model. FCoE assumes a lossless Ethernet, in which frames are never dropped (as in Fibre Channel) and that therefore does not use IP and TCP.
http://www.cisco.com/c/en/us/products/collateral/switches/nexus-5000-seriesswitches/white_paper_c11-495142.html

QUESTION 69

Hotspot

Instructions ☒

- Go through NX-OS CLI captures in Exhibits 1 through 5 to answer the questions.
- THIS TASK DOES NOT REQUIRE DEVICE CONFIGURATION.
- To access the multiple-choice questions, click the numbered boxes at the left of the top panel.
- There are four multiple-choice questions with this task. Be sure to answer all four questions before selecting the Next button.

Scenario ☒

Customer is deploying Cisco FabricPath in their new data center as shown in the topology diagram. Go through NX-OS CLI captures in Exhibits 1 through 5 to answer the questions.

Topology

Exhibit 1

```
Nexus7000-1#show feature-set
Feature Set Name        ID        State
-------------------     --------  --------
fabricpath              2         enabled
fex                     3         disabled

Nexus7000-1#
```

Exhibit 2

```
Nexus7000-1# show feature-set services fabricpath
u2rib
drap
isis_l2mp
3 services in feature set fabricpath
Nexus7000-1#
```

Exhibit 3

```
Nexus7000-1# config terminal

Nexus7000-1#(config)# fabricpath switch-id 23

Nexus7000-1#(config)#
```

Exhibit 4

```
Nexus7000-1# config terminal

Nexus7000-1#(config)# fabricpath timer allocate-delay 600

Nexus7000-1#(config)#
```

```
Exhibit 5

Nexus7000-1# config terminal
Nexus7000-1#(config)# fabricpath load-balance unicast layer3
Nexus7000-1#(config)#

Nexus7000#(config)# sh fabricpath load-balance
ECMP load-balancing configuration:
L3/L4 Preference: Mixed
Rotate amount: 14 bytes
Use VLAN: TRUE
Ftag load-balancing configuration:
Rotate amount: 3 bytes
Use VLAN: TRUE
```

On a Cisco Nexus 7000 switches what is true regarding Cisco FabricPath requirements?

A. Ensure that you have installed the Enhanced Layer 2 license and that you have installed an F Series module
B. Ensure that you have installed the Enhanced Layer 2 license and that you have installed an M Series module
C. Ensure that you have installed the Enhanced Layer 3 license and that you have installed an M Series module
D. Ensure that you have installed the Scalable Feature License license and that you have installed an F Series module

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
FabricPath switching has the following prerequisites:
You should have a working knowledge of Classical Ethernet Layer 2 functioning.
You must install the FabricPath feature set on the default and nondefault VDC before you enable FabricPath on the switch.
See Configuring Feature Set for FabricPath for information on installing the FabricPath feature set.
You are logged onto the device.

Ensure that you have installed the Enhanced Layer 2 license.
You are in the correct virtual device context (VDC). A VDC is a logical representation of a set of system resources.
You can use the switchto vdc command with a VDC number.
You are working on the F Series module.
http://www.cisco.com/en/US/docs/switches/datacenter/sw/5_x/nxos/fabricpath/configuration/guide/fp_switching.html

QUESTION 70
Hotspot

Instructions
- Go through NX-OS CLI captures in Exhibits 1 through 5 to answer the questions.
- THIS TASK DOES NOT REQUIRE DEVICE CONFIGURATION.
- To access the multiple-choice questions, click the numbered boxes at the left of the top panel.
- There are four multiple-choice questions with this task. Be sure to answer all four questions before selecting the Next button.

Scenario
Customer is deploying Cisco FabricPath in their new data center as shown in the topology diagram. Go through NX-OS CLI captures in Exhibits 1 through 5 to answer the questions.

Topology

Legend:
- Fibre Channel
- 10 Gi with FCoE
- 10 Gi
- 1 Gi

Exhibit 1

```
Nexus7000-1#show feature-set
Feature Set Name        ID        State
-------------------  --------  --------
fabricpath              2         enabled
fex                     3         disabled

Nexus7000-1#
```

Exhibit 2

```
Nexus7000-1# show feature-set services fabricpath
u2rib
drap
isis_l2mp
3 services in feature set fabricpath
Nexus7000-1#
```

Exhibit 3

```
Nexus7000-1# config terminal

Nexus7000-1#(config)# fabricpath switch-id 25

Nexus7000-1#(config)#
```

Exhibit 4

```
Nexus7000-1# config terminal

Nexus7000-1#(config)# fabricpath timer allocate-delay 600

Nexus7000-1#(config)#
```

```
Exhibit 5

Nexus7000-1# config terminal
Nexus7000-1#(config)# fabricpath load-balance unicast layer3
Nexus7000-1#(config)#

Nexus7000#(config)# sh fabricpath load-balance
ECMP load-balancing configuration:
L3/L4 Preference: Mixed
Rotate amount: 14 bytes
Use VLAN: TRUE
Ftag load-balancing configuration:
Rotate amount: 3 bytes
Use VLAN: TRUE
```

What is effect of the command "fabric path load-balance unicast layer 3"?

A. It configures F2 VDC FabricPath unicast load balancing
B. The command automatically load balances broadcast traffic
C. It configures F1/MI VDC FabricPath unicast load balancing
D. It configures M1 VDC FabricPath unicast load balancing

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The F1 cards are complemented by M1 card for routing purposes. When using M1 cards in the same virtual device context (VDC) as the F1 card, routing is offloaded to the M1 cards, and more routing capacity is added to the F1 card by putting more M1 ports into the same VDC as the F1 card.

QUESTION 71

Instructions

- Go through NX-OS CLI captures in Exhibits 1 through 5 to answer the questions.
- THIS TASK DOES NOT REQUIRE DEVICE CONFIGURATION.
- To access the multiple-choice questions, click the numbered boxes at the left of the top panel.
- There are four multiple-choice questions with this task. Be sure to answer all four questions before selecting the Next button.

Scenario

Customer is deploying Cisco FabricPath in their new data center as shown in the topology diagram. Go through NX-OS CLI captures in Exhibits 1 through 5 to answer the questions.

Topology

Legend:
- Fibre Channel
- 10 Gi with FCoE
- 10 Gi
- 1 Gi

Exhibit 1

```
Nexus7000-1#show feature-set
Feature Set Name        ID        State
--------------------    --------  --------
fabricpath              2         enabled
fex                     3         disabled

Nexus7000-1#
```

Exhibit 2

```
Nexus7000-1# show feature-set services fabricpath
u2rib
drap
isis_l2mp
3 services in feature set fabricpath
Nexus7000-1#
```

Exhibit 3

```
Nexus7000-1# config terminal

Nexus7000-1#(config)# fabricpath switch-id 23

Nexus7000-1#(config)#
```

Exhibit 4

```
Nexus7000-1# config terminal

Nexus7000-1#(config)# fabricpath timer allocate-delay 600

Nexus7000-1#(config)#
```

Exhibit 5

```
Nexus7000-1# config terminal
Nexus7000-1#(config)# fabricpath load-balance unicast layer3
Nexus7000-1#(config)#

Nexus7000#(config)# sh fabricpath load-balance
ECMP load-balancing configuration:
L3/L4 Preference: Mixed
Rotate amount: 14 bytes
Use VLAN: TRUE
Ftag load-balancing configuration:
Rotate Amount: 3 bytes
Use VLAN: TRUE
```

Customer has configured fabricpath allocate-delay to 600. What is the effect of this?

A. The allocate-delay is the time for FP to go Into forwarding state
B. It specifies the time delay for a transitioned value to be propagated throughout the network
C. It specifies the time delay for a link bringup to detect conflicts
D. The allocate-delay is the time delay for a new resource to be propagated throughout the network

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Specifies the time delay for a new resource to be propagated throughout the network.
http://www.cisco.com/web/techdoc/dc/reference/cli/nxos/commands/fpath/fabricpath_timers.html

QUESTION 72
Hotspot

Instructions

- Go through NX-OS CLI captures in Exhibits 1 through 5 to answer the questions.
- THIS TASK DOES NOT REQUIRE DEVICE CONFIGURATION.
- To access the multiple-choice questions, click the numbered boxes at the left of the top panel.
- There are four multiple-choice questions with this task. Be sure to answer all four questions before selecting the Next button.

Scenario

Customer is deploying Cisco FabricPath in their new data center as shown in the topology diagram. Go through NX-OS CLI captures in Exhibits 1 through 5 to answer the questions.

Topology

Exhibit 1

```
Nexus7000-1#show feature-set
Feature Set Name      ID      State
--------------------  -----   --------
fabricpath            2       enabled
fcx                   3       disabled

Nexus7000-1#
```

Exhibit 2

```
Nexus7000-1# show feature-set services fabricpath
u2rib
drap
isis_l2mp
3 services in feature set fabricpath
Nexus7000-1#
```

Exhibit 3

```
Nexus7000-1# config terminal

Nexus7000-1#(config)# fabricpath switch-id 23

Nexus7000-1#(config)#
```

Exhibit 4

```
Nexus7000-1# config terminal

Nexus7000-1#(config)# fabricpath timer allocate-delay 600

Nexus7000-1#(config)#
```

```
Exhibit 5

Nexus7000-1# config terminal

Nexus7000-1#(config)# fabricpath load-balance unicast layer3

Nexus7000-1#(config)#

Nexus7000#(config)# sh fabricpath load-balance

ECMP load-balancing configuration:

L3/L4 Preference: Mixed

Rotate amount: 14 bytes

Use VLAN: TRUE

Ftag load-balancing configuration:

Rotate amount: 3 bytes

Use VLAN: TRUE
```

FabricPath switch-id is 25 and load-balance is configured for L3/L4 and rotate amount is 14 byte. What information is true about FabricPath switch-id?

A. FabricPath topology requires manual configuration of switch-id which has a range from 1 to 4095
B. Every FabricPath must have a manually configured switch-id for it to form a FabricPath topology
C. FabricPath topology requires manual configuration of switch-id which has a range from 1 to 4099
D. You do not have to manually assign a switch ID unless you are running a virtual port channel plus (vPC*) because the system assigns a switch ID for you when you enable FabricPath

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
fabricpath switch-id (vPC)
To configure a virtual port channel plus (vPC+) switch ID, use the fabricpath switch-id command. To remove the FabricPath switch from a vPC domain, use the no form of this command.
fabricpath switch-id switch-id

no fabricpath switch-id [switch-id]
Usage Guidelines
You do not have to manually assign a switch ID (unless you are running a vPC+); the system assigns a switch ID for you when you enable FabricPath.
Note You must assign the same vPC+ switch ID to each of the two vPC+ peer devices before they can form an adjacency.
This command requires an Enhanced Layer 2 license.
Examples
This example shows how to configure a vPC+ switch ID on a FabricPath-enabled device:
switch# configure terminal switch(config)# vpc domain 1
switch(config-vpc-domain)# fabricpath switch-id 1
Configuring fabricpath switch id will flap vPCs. Continue (yes/no)? [no]

QUESTION 73
Hotspot

Instructions

- THIS TASK DOES NOT REQUIRE DEVICE CONFIGURATION.
- Click Cisco Nexus 5548 to gain console access. No console or enable passwords are required.
- To access the multiple-choice questions, click the numbered boxes on the left of the top panel.
- There are four multiple-choice questions with this task.

Scenario

Customer is deploying Cisco Nexus 5548 switch with FCoE in their new data center, as shown in the topology diagram. Click Nexus5548 icon to run show commands and answer the questions.

Topology

LUN   MDS 9124   Nexus 2248TP

Nexus 5548

UCS C200

Legend:

▬▬▬▬	Fibre Channel
▬ ▬ ▬	10 Gi with FCoE
▬▬▬▬	10 Gi
··········	1 Gi

Nexus 5548

Nexus 5548#

What is the status of FCoE license on Cisco Nexus 5548 switch?

A. FCoE license is not installed
B. FCoE license is installed, but it is expired
C. FCoE license is installed and status is enabled
D. FCoE license does not need to be installed because it is part of ENTERPRISE_PKG

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 74
Hotspot

Instructions ☒

- THIS TASK DOES NOT REQUIRE DEVICE CONFIGURATION.
- Click Cisco Nexus 5548 to gain console access. No console or enable passwords are required.
- To access the multiple-choice questions, click the numbered boxes on the left of the top panel.
- There are four multiple-choice questions with this task.

Scenario ☒

Customer is deploying Cisco Nexus 5548 switch with FCoE in their new data center, as shown in the topology diagram. Click Nexus5548 icon to run show commands and answer the questions.

Topology

LUN   MDS 9124   Nexus 2248TP

Nexus 5548   UCS C200

Legend:

Fibre Channel	
10 Gi with FCoE	
10 Gi	
1 Gi	

Topology

LUN   MDS 9124   Nexus 2248TP

Nexus 5548   UCS C200

Legend:

▬▬▬▬	Fibre Channel
─ ─ ─	10 Gi with FCoE
───────	10 Gi
··········	1 Gi

Ethernet interface 1/5 on Cisco Nexus 5548 is connected to Cisco UCB C220 rack server. What is the status of Ethernet 1/5 interface for FCoE functionality?

A. Interface reset on Ethernet 1/5 is preventing the FCoE connection from coming up
B. MTU size of 1500 on Ethernet interface 1/5 needs to be changed for FCoE to come UP
C. Cisco Nexus 5548 needs a layer 3 daughter card for FCoE to come UP on the Ethernet interface 1/5
D. Ethernet interface 1/5 is operational for FCoE and the status is UP

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 75
Hotspot

Instructions ☒

- THIS TASK DOES NOT REQUIRE DEVICE CONFIGURATION.
- Click Cisco Nexus 5548 to gain console access. No console or enable passwords are required.
- To access the multiple-choice questions, click the numbered boxes on the left of the top panel.
- There are four multiple-choice questions with this task.

Scenario ☒

Customer is deploying Cisco Nexus 5548 switch with FCoE in their new data center, as shown in the topology diagram. Click Nexus5548 icon to run show commands and answer the questions.

Topology ☒

LUN

MDS 9124

Nexus 2248TP

Nexus 5548

UCS C200

Legend:

▬▬▬	Fibre Channel
– – –	10 Gi with FCoE
▬▬▬	10 Gi
··········	1 Gi

Nexus5548

Nexus5548#
Nexus5548#
Nexus5548#
Nexus5548#

What is the status of FC interface associated with ethernet 1/5 indicate?

A. Trunk VSAN 11 is isolated
B. Interface vfc 5 is up and running for the assigned VSAN
C. Trunk VSAN 11 is initializing
D. VSAN to FC mapping is not working as expected

Correct Answer: B
Section: (none)
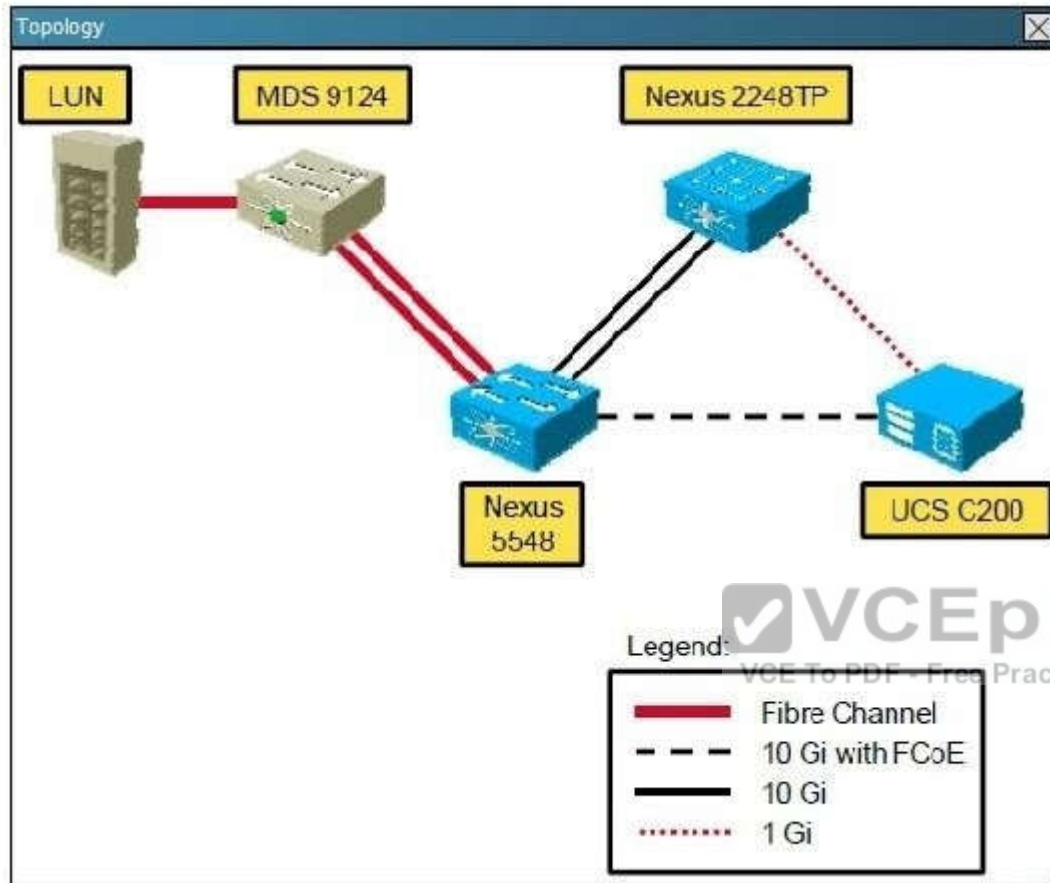**Explanation**

Explanation/Reference:

QUESTION 76
Hotspot

Instructions ☒

- THIS TASK DOES NOT REQUIRE DEVICE CONFIGURATION.
- Click Cisco Nexus 5548 to gain console access. No console or enable passwords are required.
- To access the multiple-choice questions, click the numbered boxes on the left of the top panel.
- There are four multiple-choice questions with this task.

Scenario ☒

Customer is deploying Cisco Nexus 5548 switch with FCoE in their new data center, as shown in the topology diagram. Click Nexus5548 icon to run show commands and answer the questions.

Topology

LUN   MDS 9124   Nexus 2248TP

Nexus 5548

UCS C200

Legend:

▬▬▬	Fibre Channel
─ ─ ─	10 Gi with FCoE
───	10 Gi
·········	1 Gi

```
Nexus5548


Nexus5548#
Nexus5548#
Nexus5548#
Nexus5548#
```

When configuring FCoE VLANs and Virtual Fiber Channel (vFC) interfaces, what guidelines must be followed must be followed?

A. Each vFC interface must be bound to an FCoE-enabled Ethernet or EtherChannel interface or to the MAC address of a remotely connected adapter
B. Each FC Interface must be bound to an FCoE-enabled Ethernet or EtherChannel interface or to the MAC address of a remotely connected adapter
C. Each vFC Interface must be bound to an FC enabled Ethernet or EtherChannel interface or to the MAC address of a remotely connected adapter
D. Each vFC Interface must be bound to an FCoE-enabled vFC or EtherChannel interface or to the MAC address of a remotely connected adapter

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 77

Which three options are capabilities of the Cisco Nexus 7000 Series Switch? (Choose three.)

A. All interface and supervisor modules are accessible from the front.
B. All interface and supervisor modules are accessible from the rear.
C. single power supply only
D. multiple power supply option for redundancy
E. up to 180.7 Tbps forwarding capacity with Fabric-2 modules with 10-slot switches
F. up to 18.7 Tbps forwarding capacity with Fabric-2 modules with 18-slot switches

Correct Answer: ADF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 78
Which four options are capabilities of the Cisco Nexus 5000 and 5500 Series Switch? (Choose four.)

A. line rate
B. managed by a parent switch
C. lossless 10 Gigabit Ethernet
D. lossless 100 Gigabit Ethernet
E. low latency
F. extremely low latency
G. hosts a virtual supervisor module

Correct Answer: ACEG
Section: (none)
Explanation

Explanation/Reference:

QUESTION 79
Which three options are capabilities of the Cisco Nexus 7000 Series Supervisor Module? (Choose three.)

A. hardware forwarding on the supervisor module

B. fully decoupled control plane and data plane with no forwarding on the supervisor module

C. Sup2 requires Cisco NX-OS 5.1 or later.

D. Sup2 requires Cisco NX-OS 6.1 or later.

E. Sup2E supports 8+1 VDC with the N7K-VDC1K9 license per chassis.

F. Sup2 supports 8+1 VDCs with the N7K-VDC1K9 license per chassis.

Correct Answer: BDE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 80
Which Cisco NX-OS feature allows transparent Layer 2 extension between sites?

A. FabricPath

B. ETV

C. OTV

D. vPC

E. LISP

F. TrustSec

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 81
Drag and Drop
Drag the network characteristics on the left to the most appropriate design layer on the right.

Select and Place: