



- Expert Verified, Online, **Free**.

What is an advantage of using OTV as compared to VPLS for data center redundancy?

- A. prevents loops on point-to-point links
- B. provides head-end replication
- C. uses a proactive MAC advertisement
- D. provides full-mesh connectivity

Correct Answer: A

Reference:

<https://community.cisco.com/t5/data-center-documents/understanding-overlay-transport-virtualization-otv/ta-p/3151502#toc-hid-1043251551>

Community vote distribution

C (100%)

🗨️ **k3rnelpanicpj** 1 month, 2 weeks ago

Selected Answer: C

Answer C is correct.

Why:

VPLS also prevent loops

otv avoid head-end replication

vpls also use full-mesh connectivity

upvoted 1 times

🗨️ **vtache** 1 year, 8 months ago

I would say answer is B

upvoted 1 times

🗨️ **TroyNguyen** 1 year, 10 months ago

Hi all,

I didn't pass after 3 times. Poor me. Do any one here can help/share right answer validated for me to get out of that ? Thank you so much.

upvoted 1 times

🗨️ **husam_shabeeb** 1 year, 10 months ago

I didnt pass , new 20-25 questions >19-4-2023 total 61 3-4 drag and drop

upvoted 1 times

🗨️ **ronnius** 2 years, 4 months ago

Yes, it is C

https://www.cisco.com/c/dam/en_us/solutions/industries/docs/gov/otv.pdf

"Flooding Based Learning vs Control-Plane Based Learning

Move to a Control Plane protocol that proactively advertises MAC

addresses and their reachability instead of the current flooding

mechanism."

upvoted 4 times

🗨️ **sadiq990** 2 years, 10 months ago

Im sorry answer is A

upvoted 1 times

🗨️ **sadiq990** 2 years, 10 months ago

Overlay transportation introduces the concept of "MAC routing," which means a control plane protocol is used to exchange MAC reachability information between network devices providing LAN extension functionality. This is a significant shift from Layer 2 switching that traditionally leverages data plane learning, and it is justified by the need to limit flooding of Layer 2 traffic across the transport infrastructure. As outlined in this chapter, Layer 2 communication between sites resembles routing more than switching. If the destination MAC address information is unknown, traffic is dropped (not flooded), preventing the waste of precious bandwidth across the WAN

Answer is C

upvoted 1 times

🗨️ **Manu_75** 2 years, 11 months ago

Correct answer is C -- see slide 18:

<https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2018/pdf/BRKDCN-2931.pdf>

upvoted 1 times

🗨️ **gu33mis** 3 years, 1 month ago

Shouldn't this be option - B?

OTV supports head-end replication when multicast is not supported in the core, so no need to broadcast the traffic to all sites.

upvoted 1 times

🗨️ **MajorSalieri** 3 years ago

No, Head-end replication is only supported by VXLAN through static mappings of VTEPs. A is correct

upvoted 1 times

🗨️ **lurker8000** 3 months, 2 weeks ago

Loop on p2p? I don't think so...

upvoted 1 times

🗨️ **dzilla315** 4 years ago

FYI after taking the test and most of the questions are not on this webpage. Maybe 5-10. best of luck

upvoted 4 times

DRAG DROP -

A failure occurs on the network between two BFD and OSPF neighbors. Drag and drop the protocol actions from the left into the correct order on the right.

Select and Place:

BFD notifies the local OSPF process that the BFD neighbor is no longer reachable.	step 1
If an alternate path is available, routers immediately start converging on it.	Step 2
The BFD neighboring session with the OSPF peer is terminated.	Step 3
The local OSPF process tears down the OSPF neighbor relationship.	Step 4

Correct Answer:

BFD notifies the local OSPF process that the BFD neighbor is no longer reachable.	The BFD neighboring session with the OSPF peer is terminated.
If an alternate path is available, routers immediately start converging on it.	BFD notifies the local OSPF process that the BFD neighbor is no longer reachable.
The BFD neighboring session with the OSPF peer is terminated.	The local OSPF process tears down the OSPF neighbor relationship.
The local OSPF process tears down the OSPF neighbor relationship.	If an alternate path is available, routers immediately start converging on it.

Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fs_bfd.html#wp1238898

 **msagimlar** 4 months ago

Yes, correct

upvoted 2 times

Which multicast rendezvous point redundancy mode is valid for Bidirectional PIM?

- A. Embedded RP
- B. Phantom RP
- C. MSDP
- D. PIM anycast RP

Correct Answer: D

Reference:

https://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/ip-multicast/whitepaper_c11-508498.html

Community vote distribution

B (100%)

 **CHEdot** Highly Voted 4 years, 6 months ago

Answer must be B.

Check the Phantom RP section in:

https://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/ip-multicast/whitepaper_c11-508498.html

upvoted 9 times

 **Mr_Cipher** 4 years, 5 months ago

I reckon Anycast-RP is the correct answer, refer back to your doc, you'll see the Anycast-RP provides RP failover and redundancy.

upvoted 1 times

 **FomaLermont85** 4 years, 3 months ago

See page 8 in https://www.cisco.com/c/dam/en/us/products/collateral/ios-nx-os-software/multicast-enterprise/prod_white_paper0900aecd80310db2.pdf. It says "Because multicast source information is no longer available in Bidir, the Anycast/MSDP mechanism used to provide redundancy in sparse-mode is not an option for Bidir". So, correct answer should be B, IMHO

upvoted 7 times

 **Font** Highly Voted 3 years, 10 months ago

should be B :

<https://community.cisco.com/t5/networking-documents/rp-redundancy-with-pim-bidir-phantom-rp/ta-p/3117191>

upvoted 5 times

 **Wasamela** Most Recent 8 months, 1 week ago

Selected Answer: B

Agree with CHEdot the correct answer is B as the Phantom RP provides dynamic failover in case the primary RP becomes unavailable.

upvoted 1 times

 **DGriff** 9 months ago

A few hints:

MSDP is IPV4 only and used to blamance PIM registrations

PIM anycast RP substitute and extents the role of MSDP to RPs sets and supports IPv6

Phantom RP is specifically for PIM-BiDir, the nature of this model is birectional flowing between source/recievers, and Phantom RPs are just transit points along the flow path (no registration role)

upvoted 1 times

 **mojzucha** 3 years, 4 months ago

It is Phantom RP of course

upvoted 1 times

An engineer deploys LISP VM mobility. Which feature is configured on the interfaces that have VM mobility enabled?

- A. IP redirects
- B. flow control
- C. proxy ARP
- D. HSRP

Correct Answer: C

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nx-os/lisp/configuration/guide/b_NX-OS_LISP_Configuration_Guide/b_NX-OS_LISP_Configuration_Guide_chapter_010.html

Community vote distribution

C (100%)


 **Fcpoultry** 8 months ago

Selected Answer: C

LISP VM mobility across subnets requires that the same MAC address is configured across all HSRP groups that allow dynamic EIDs to roam. You must enable the Proxy Address Resolution Protocol (ARP) for the interfaces that have VM mobility enabled across subnets.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nx-os/lisp/configuration/guide/b_NX-OS_LISP_Configuration_Guide/b_NX-OS_LISP_Configuration_Guide_chapter_010.html

upvoted 1 times

 **jack98** 2 years, 9 months ago

correct

upvoted 1 times

 **Magneto** 3 years ago

Correct

upvoted 1 times

What are two advantages of using Cisco vPC over traditional access layer designs? (Choose two.)

- A. supports Layer 3 port channels
- B. disables spanning-tree
- C. no spanning-tree blocked ports
- D. uses all available uplink bandwidth
- E. maintains single control plane

Correct Answer: *CD*

Reference:

https://www.cisco.com/c/dam/en/us/td/docs/switches/datacenter/sw/design/vpc_design/vpc_best_practices_design_guide.pdf

🗨️ 👤 **AboudMokh** 7 months ago

correct

upvoted 1 times

🗨️ 👤 **Magneto** 1 year ago

vPC provides the following technical benefits:

- Eliminates Spanning Tree Protocol (STP) blocked ports
- Uses all available uplink bandwidth
- Allows dual-homed servers to operate in active-active mode
- Provides fast convergence upon link or device failure
- Offers dual active/active default gateways for servers vPC also leverages native split horizon/loop management provided by port-channeling technology: a packet entering a port-channel cannot immediately exit that same port-channel.

upvoted 2 times

DRAG DROP -

Drag the appropriate from the left onto the current UDLD modes of operation on the right.

Select and Place:

Uses STP if UDLD information times out.

Detects when a port is stuck.

Disables a port that has a high error rate.

Allows a port that has a high error rate.

Normal Mode

Aggressive Mode

Correct Answer:

Uses STP if UDLD information times out.

Detects when a port is stuck.

Disables a port that has a high error rate.

Allows a port that has a high error rate.

Normal Mode

Uses STP if UDLD information times out.

Allows a port that has a high error rate.


Aggressive Mode

Detects when a port is stuck.

Disables a port that has a high error rate.

Reference:

<https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/10591-77.html>

 **Magneto** 6 months, 3 weeks ago

Correct

upvoted 1 times


What is a design consideration when implementing FSPF?

- A. Routes are based on the domain ID.
- B. Routes are based on the distance vector protocol.
- C. FSPF runs only on F Ports.
- D. FSPF runs on a per-chassis basis.

Correct Answer: A

Community vote distribution


A (100%)

 **vtache** 8 months, 1 week ago

Selected Answer: A

i would say the answer is A

upvoted 1 times

 **sadiq990** 1 year, 10 months ago

Selected Answer: A

Domain ID: The domain ID is assigned to each FC switch in a fabric and is unique to that switch. Because the upper FCIDs are reserved for various FC fabric services, such as 0xFFFFF for broadcast and 0xFFFFC for name server, the domain ID field restricts the maximum number of FC switches that can be connected in a fabric or VSAN to 239. Routing decisions are made based on domain ID bits of FCID

Answer is A

upvoted 2 times

Which two features are provided by deploying an OOB management network in a Cisco Nexus data center? (Choose two.)

- A. Layer 3 path for monitoring purposes
- B. Layer 2 path for server traffic
- C. Layer 2 path for a vPC peer link
- D. Layer 3 path for vPC keepalive packets
- E. Layer 3 path for server traffic

Correct Answer: BC

Reference:

<https://www.cisco.com/c/dam/en/us/td/docs/solutions/CVD/Aug2014/CVD-DataCenterDesignGuide-AUG14.pdf>


Community vote distribution

AD (100%)

 **CHEdot** Highly Voted 4 years, 6 months ago


Should be AD?

upvoted 14 times

 **ZsF** 4 years, 3 months ago

A for sure, but.. It provides L2 for the keepalive messages.. sooo.. no correct answer..

upvoted 2 times

 **petarbr** 2 years, 1 month ago

Not correct, OOB does not anything regarding VPC. Indirectly, excluding B,C and E, A and D stays.

upvoted 1 times

 **abdo982** Highly Voted 4 years, 5 months ago

i think AD is correct

upvoted 6 times

 **missionflamingo** Most Recent 1 month, 1 week ago

Selected Answer: AD

mgmt0 is a OOB L3 interface for control plain traffic only

upvoted 1 times

 **Wasamela** 8 months ago

Selected Answer: AD

Answer should be A and D

upvoted 1 times

 **NeoTheOne** 2 years, 10 months ago

AD is the correct answer

upvoted 1 times

 **alexshang** 3 years, 1 month ago

AD. OOB is not for server traffic

upvoted 1 times

 **mojzucha** 3 years, 4 months ago

I vote for A and D. As said earlier, OOB is not for server traffic nor for peer link.


upvoted 1 times

 **FonT** 3 years, 10 months ago

OOB is not for server traffic and not for peer link.



so it's D and A

upvoted 3 times

 **Magneto** 3 years, 6 months ago

You are right

upvoted 2 times

  **Whirly33** 4 years ago

You wouldn't design for server traffic over the management interface. Therefore, all the answers with server traffic are out. Cisco also only recommends the VPC Keepalive feature be designated to its own link or via the management interface. It does not recommend the same for the Peer feature. Therefore, the only valid answers are the for monitoring and/or VPC Keepalive.

upvoted 1 times

  **dzilla315** 4 years, 1 month ago

its B C..



upvoted 1 times

The management of the Cisco Nexus switches is provided over an isolated out-of-band network. The VDC feature is configured on the Cisco Nexus core switches.

How is out-of-band management access provided for each VDC?

- A. All the VDC have the same out-of-band IP address.
- B. Each VDC has a dedicated out-of-band Ethernet management port.
- C. Each VDC has a unique out-of-band IP address from the same IP subnet.
- D. Each VDC has a unique out-of-band IP address from different IP subnets among VDCs.

Correct Answer: C

  **sadiq990** 4 months, 2 weeks ago

Virtual Device Contexts (VDC) allow you to carve out multiple virtual switches from a single physical Nexus switch. Each VDC is logically isolated from every other VDC on a switch. Therefore just like with physical switches, in order to trunk or route traffic between them, physical interfaces and cabling is required to connect two or more VDCs together before this can happen. A useful VDC feature is that the mgmt0 interface on the Supervisor can be assigned a different IP address by each VDC.

VDC management can be done via virtualizing the management port so that, each VDC will get logical instance of management port and where we have to provide the separate IP address to each Virtual management Port present in VDC from the same subnet

Answer is C

upvoted 2 times

Which two configurations allow for routing traffic between two VDCs? (Choose two.)

- A. Connect the VDCs to an external Layer 3 device.
- B. Cross-connect the ports between the VDCs.
- C. Create VRF-aware software infrastructure interfaces.
- D. Create a policy map in the default VDC that routes traffic between the VDCs.
- E. Create interfaces in the VDC that can be accessed by another VDC.

Correct Answer: AD

Community vote distribution

AB (100%)

 **CHEdot** Highly Voted 4 years, 6 months ago

Should be A&B

upvoted 21 times

 **Wasamela** Most Recent 8 months ago

Selected Answer: AB

A&B should be the right answer.

upvoted 1 times

 **mojzucha** 3 years, 4 months ago

A and B of course

upvoted 1 times

 **aelmsieh** 3 years, 9 months ago

The Cisco NX-OS software does not support direct communication between VDCs on a single physical device. You must make a physical connection from a port allocated to one VDC to a port allocated to the other VDC to allow the VDCs to communicate.

so A&B

upvoted 4 times

Which protocol prevents fiber failures from causing a Layer 2 loop in a dark fiber DCI?

- A. BFD
- B. UDLD
- C. MCP
- D. LLDP

Correct Answer: B

Community vote distribution

C (50%)

B (50%)

🗨️ 👤 **Wasamela** 8 months, 1 week ago

Selected Answer: B

UDLD is specifically designed to detect unidirectional link failures on fiber optic connections.

upvoted 1 times

🗨️ 👤 **NeoTheOne** 2 years, 10 months ago

Answer is UDLD

DCI: DataCenter Interconnect (not ACI - which is another Cisco product)

Dark Fiber: is the marketing term for a direct fiber cable dug in the ground between two DCs.

UDLD is Uni-directional Link Detection protocol that is used to detect if there is a problem with rx/tx in the fiber strands that may cause switching loop.

upvoted 2 times

🗨️ 👤 **sadiq990** 2 years, 10 months ago

Im sorry answer is correct Udld

upvoted 1 times

🗨️ 👤 **sadiq990** 2 years, 10 months ago

Selected Answer: C

The answer is wrong.

The ACI fabric does not participate in the Spanning Tree Protocol. Instead, it implements the Mis-Cabling Protocol (MCP) to detect loops. MCP works in a complementary manner with STP that is running on external Layer 2 networks and handles bridge protocol data unit (BPDU) packets that access ports receive.

The answer is MCP

upvoted 1 times

DRAG DROP -


Drag and drop the data center technologies from the left onto the correct descriptions on the right. Not all technologies are used.

Select and Place:

cut-through switching	connects a classical Ethernet vPC domain and a Cisco FabricPath cloud to interoperate
OTV	extends Layer 2 domains across distributed data centers
store-and-forward switching	performs Layer 2 lookup as soon as the destination MAC is received
VPC	two separate control planes, performs device aggregation
VPC+	

Correct Answer:

cut-through switching	VPC+
OTV	OTV
store-and-forward switching	cut-through switching
VPC	VPC
VPC+	

 **Andcak33** 3 months, 1 week ago

Correct!

upvoted 2 times

What are two characteristics of VRF-Lite on the Cisco NX-OS platform? (Choose two.)

- A. Layer 3 interfaces can belong to multiple VRFs.
- B. VRF-Lite interfaces can be Layer 2.
- C. Overlapping IP addresses are permitted on different VRFs.
- D. Interfaces that belong to VRFs can be physical or logical.
- E. VRF-Lite interfaces support MPLS.

Correct Answer: CD

Community vote distribution

CD (100%)

🗨️ 👤 **sadiq990** 4 months, 2 weeks ago

Selected Answer: CD

VRF can be either physical, such as Ethernet ports, or logical, such as VLAN SVIs, but a Layer 3 interface cannot belong to more than one VRF at any time.

With the VRF-lite feature, the switch supports multiple VRFs in customer edge (CE) switches. VRF-lite allows a service provider to support two or more Virtual Private Networks (VPNs) with overlapping IP addresses using one interface.

upvoted 3 times

🗨️ 👤 **FomaLermont85** 1 year, 9 months ago

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/unicast/521_N11/cisco_n5k_layer3_ucast_cfg_rel_521_N1_1/13_virtual.html

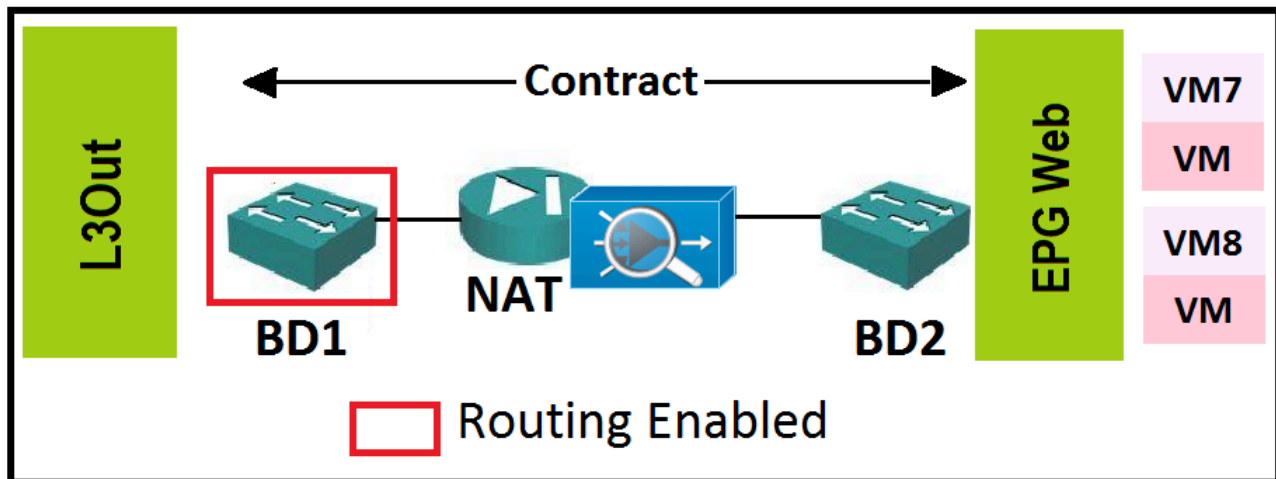
upvoted 2 times

🗨️ 👤 **Magneto** 1 year ago

Correct

upvoted 2 times

Refer to the exhibit.



What must be the default gateway of the servers in EPG Web?

- A. subnet address of BD1
- B. subnet address of L3Out
- C. subnet address of BD2
- D. NAT appliance

Correct Answer: D

lurker8000 3 months, 3 weeks ago

Why not L3Out, just curious? The NAT device may have to translate the destination, but it's also possible to place the GW on L3Out is it now?
upvoted 1 times

Woodie 2 years, 4 months ago

Because of the NAT, the server subnet and and bd2 subnet are different, so the servers cannot use bd2 subnet gateway. Therefore only the NAT device knows the local gateway for the servers. So I think answer D is correct.
upvoted 2 times

NeoTheOne 2 years, 9 months ago

D is correct because routing is not enabled on the BridgeDomain so the NAT appliance is the gateway for that bridge domain.
upvoted 1 times

Hanbur 3 years, 1 month ago

Hi all,

The answer is A.

Ref:<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-734298.html>
see Figure 35,The L4-L7 device provides the default gateway for the servers on BD1.

upvoted 1 times

AboudMokh 3 years, 1 month ago

Wrong, It clearly says "The L4-L7 device provides the default gateway for the servers." which in this case the NAT device not the BD1. It also explains "The subnet address of the outside bridge domain is the default gateway of the service appliance."

So right Answer is D

upvoted 2 times

fayssal 3 years, 5 months ago

Hi, I think that answer is A



upvoted 1 times

gensek87 4 years, 2 months ago

Routed Mode with L3Out and NAT:

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-734298.html>

upvoted 1 times

  **gensek87** 4 years, 2 months ago

looks like you need more information in this question(deployment mode) to make the right choice...



upvoted 1 times

  **abdo982** 4 years, 5 months ago

i think it should be C

the gateway of the EPG is the subnet address of the BD

upvoted 2 times

  **LionNS** 4 years, 4 months ago

Hello abdo982. It cannot be C because the device shown in the graph for this answer is a Layer 2 Switch (Layer 3 switch has a different symbol).

upvoted 2 times

Which element is the main functional component of the architecture that separates the internal fabric and the external network that connects sites when VXLAN EVPN is deployed multisite?

- A. border leaf
- B. service leaf
- C. border gateway
- D. service gateway

Correct Answer: C

Reference:

https://dcloud-cms.cisco.com/demo_news/cisco-dcnm-11-3-for-vxlan-evpn-multi-site-deployments-lab-v1

Community vote distribution

C (100%)

🗨️ 👤 **husam_shabeeb** 5 months ago

Ethernet VPN (EVPN) which is used as the overlay control plane and provides virtual connectivity between different layer 2/3 domains over an IP or MPLS network. Virtual extensible LANs (VXLAN), a common network virtualization overlay protocol that expands the layer 2 network address space from 4,000 to 16 million.

upvoted 2 times

🗨️ 👤 **sadiq990** 1 year, 4 months ago

Selected Answer: C

The main functional component of the EVPN Multi-Site architecture is the border gateway, or BGW. BGWs separate the fabric-side (site-internal fabric) from the network that interconnects the sites (site-external DCI) and mask the site-internal VTEPs.

link:<https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/white-paper-c11-739942.html>

upvoted 1 times

Which feature must be configured to connect a classical Ethernet network to a data center network so that the data center network appears to be one large switch.

- A. VPC+
- B. UDLD
- C. EVPN
- D. OSPF

Correct Answer: C

Community vote distribution

A (50%) C (50%)

Mr_Cipher **Highly Voted** 4 years, 5 months ago

It should be A. VPC+, i guess
upvoted 11 times

bentoj 4 years, 1 month ago

It suggest that we are talking about FabricPath => vPC+
upvoted 4 times

Wasamela **Most Recent** 8 months, 1 week ago

Selected Answer: C

Answer EVPN, as it enables learning and forwarding paths across the VXLAN fabric, hence creating a logical Layer 2 network that appears as a single large switch.
upvoted 1 times

husam_shabeeb 1 year, 11 months ago

Ethernet VPN (EVPN) which is used as the overlay control plane and provides virtual connectivity between different layer 2/3 domains over an IP or MPLS network. Virtual extensible LANs (VXLAN), a common network virtualization overlay protocol that expands the layer 2 network address space from 4,000 to 16 million.
upvoted 1 times

husam_shabeeb 1 year, 11 months ago

What is vPC and vPC+?
Image result for VPC+

vPC connects hosts and other devices to the network. It does this in a highly available way, using many active links. vPC+ is an extension to traditional vPC. It connects Classical Ethernet devices to a Fabricpath domain.

vPC Plus - Network Direction

networkdirection.net

<https://networkdirection.net/virtual-port-channels-vpc>

Search for: What is vPC and vPC+?

What is vPC in Cisco?

Image result for VPC+

A virtual port channel (vPC) allows links that are physically connected to two different Cisco Nexus 7000 or 9000 Series devices to appear as a single port channel by a third device. The third device can be a switch, server, or any other networking device that supports port channels. Nov 9, 2022

upvoted 1 times

NeoTheOne 2 years, 10 months ago

The correct answer is FabricPath (which is not an option) But VPC+ is VPC with Fabric path so A is most likely right answer.
upvoted 2 times

sadiq990 2 years, 10 months ago

Selected Answer: A

A vPC allows links that are physically connected to two different Nexus switches to appear as single portchannel to a server or another switch.

upvoted 1 times

  **mojzucha** 3 years, 4 months ago

VPC+ is the right option

upvoted 3 times

When planning to deploy a load-balancing service as a managed node, which ACI configuration must be included?

- A. BGP ECMP
- B. Layer 4 to Layer 7 device package
- C. inline load-balancing device
- D. ITD

Correct Answer: *B*

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/ACI_Best_Practices/b_ACI_Best_Practices/b_ACI_Best_Practices_chapter_0110.html

Currently there are no comments in this discussion, be the first to comment!

What are two reasons to select OTV as the DCI solution to connect multisite topologies? (Choose two.)

- A. It propagates hosts reachability without support of traffic flooding.
- B. Layer 3 failures do not propagate beyond the OTV edge device.
- C. It extends the spanning tree between data centers.
- D. It is an open standard.
- E. It constrains HSRP hello messages to each data center.

Correct Answer: BC

Reference:

https://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Data_Center/DCI/whitepaper/DCI3_OTV_Intro/DCI_1.html

Community vote distribution

AE (67%)

BE (33%)

 **CHEdot** Highly Voted 4 years, 6 months ago

Probably A&B.

C is wrong. OTV do not extend spanning-tree.

See "Where Is My Spanning-Tree Root with OTV?" on page 5

https://www.cisco.com/c/dam/en/us/products/collateral/switches/nexus-7000-series-switches/guide_c07-728315.pdf

upvoted 17 times

 **Magneto** 3 years, 6 months ago

Correct

upvoted 1 times

 **Font** Highly Voted 3 years, 10 months ago


I think A E

https://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Data_Center/DCI/whitepaper/DCI3_OTV_Intro/DCI_1.html

OTV introduces the concept of "MAC routing," [...] it is justified by the need to limit flooding of Layer 2 traffic across the transport infrastructure.

The last capability introduced by OTV is to filter First Hop Redundancy Protocol (FHRP—HSRP, VRRP, and so on) messages across the logical overlay.

upvoted 9 times

 **Luc1f3R** 3 years, 5 months ago

Why is this one not upvoted much?

<https://community.cisco.com/t5/data-center-documents/cisco-otv-and-fhrp-isolation/ta-p/3136485>


upvoted 1 times

 **CoAsT_x** Most Recent 2 months ago

Selected Answer: AE

I definitely believe A and E as well. If you read the link in the reference whitepaper in the suggested answer above, you can literally find A. E. stated in there.


upvoted 1 times

 **Wasamela** 8 months, 1 week ago

Selected Answer: BE

My bad, B&E

upvoted 1 times

 **Wasamela** 8 months, 1 week ago

Selected Answer: AE

A&E are correct.

upvoted 1 times

🗨️ 👤 **mojzucha** 3 years, 4 months ago

A and E are the options of your choice
upvoted 5 times

🗨️ 👤 **LionNS** 4 years, 4 months ago

In the URL provided as reference says: "This is achieved by leveraging the same control plane protocol used for the exchange of MAC address information, without the need of extending the Spanning-Tree Protocol (STP) across the overlay".

So as CHEdot commented, C is incorrect.
upvoted 2 times

🗨️ 👤 **LionNS** 4 years, 4 months ago

I believe B and E are correct. I chose E because "The last capability introduced by OTV is to filter First Hop Redundancy Protocol (FHRP—HSRP, VRRP, and so on) messages across the logical overlay."
upvoted 3 times



Which technology enables Layer 2 extension between remote data center sites?

- A. FCIP
- B. GRE
- C. DMVPN
- D. OTV

Correct Answer: *D*

Reference:

<https://netcraftsmen.com/spanning-layer-2-between-data-centers/>

  **Andcak33** 3 months, 1 week ago

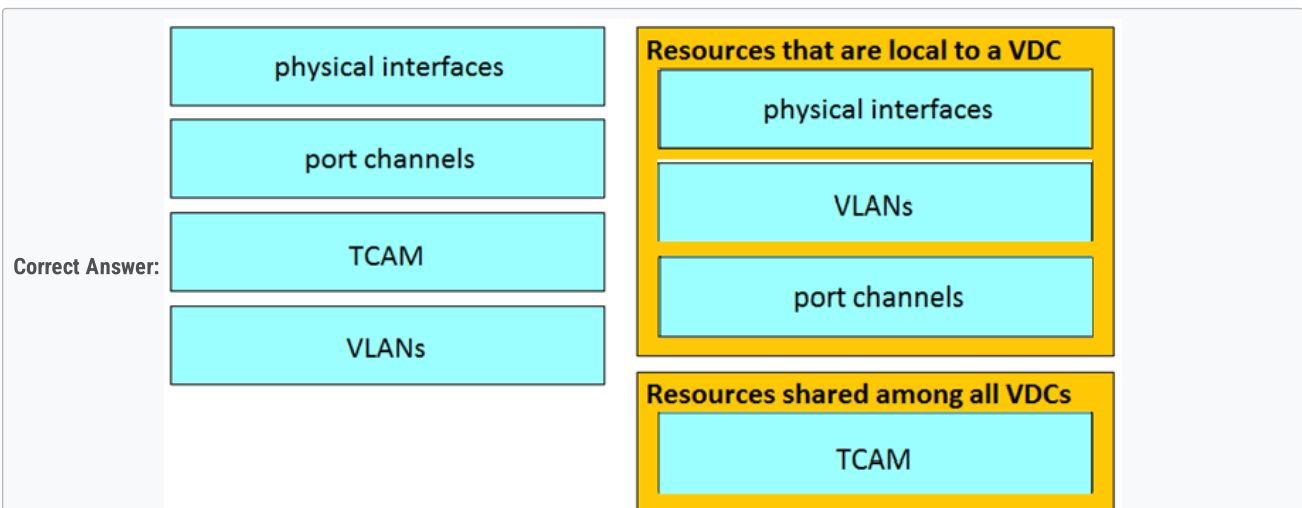
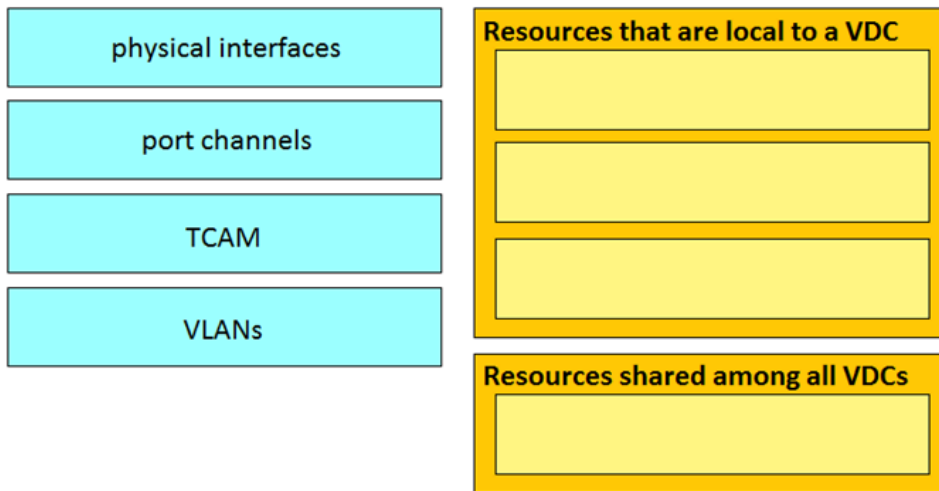
Correct!

upvoted 1 times

DRAG DROP -

Drag and drop the Cisco Nexus 7000 Series Switch resources from the left onto the correct groups on the right.

Select and Place:



Reference:

https://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-Virtual-Device-Context-Configuration-Guide/creating-vdc.html

Font Highly Voted 2 years, 4 months ago

I confirm TCAM is shared

"The FIB TCAM entries are system wide resources that are shared across virtual device contexts (VDC) configured on the module. Table 16-2 describes the supported maximum FIB scale entries on the Nexus 7000 system configuration per route-type."

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nx-os/unicast/configuration/guide/b-7k-Cisco-Nexus-7000-Series-NX-OS-Unicast-Routing-Configuration-Guide-Release/n7k_unicast_managing_rib_fib.html

upvoted 5 times

Font Most Recent 8 months, 1 week ago

"The FIB TCAM entries are system wide resources that are shared across virtual device contexts (VDC) configured on the module. Table 16-2 describes the supported maximum FIB scale entries on the Nexus 7000 system configuration per route-type."

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nx-os/unicast/configuration/guide/b-7k-Cisco-Nexus-7000-Series-NX-OS-Unicast-Routing-Configuration-Guide-Release/n7k_unicast_managing_rib_fib.html

upvoted 1 times

Andcak33 9 months, 2 weeks ago

TCAM sections are dedicated for each VDC, TCAM as a whole is not.

upvoted 1 times

🗨️ 👤 **Woodie** 10 months, 3 weeks ago

Physical Interfaces, Port-Channels and VLANs are not shared between VDCs. So the most reasonable option for a shared resource is the TCAM.

upvoted 2 times

🗨️ 👤 **bittenbt** 1 year, 1 month ago

TCAM is dedicated and local to a VDC. P64 in the course book.

upvoted 1 times

🗨️ 👤 **msagimlar** 2 years, 1 month ago

True answer ?

upvoted 1 times

🗨️ 👤 **bentoj** 2 years, 7 months ago

TCAM is not shared, it's dedicated to each VDC

upvoted 1 times

DRAG DROP -

Drag and drop the feature descriptions from the left onto the correct UCS Fabric Interconnect modes on the right.

Select and Place:

Links are active-active.	End-Host Mode
MAC address learning occurs on the uplink ports.	
STP runs on the uplink ports.	Switch Mode
This is the default mode of operation.	

Correct Answer:

Links are active-active.	End-Host Mode
MAC address learning occurs on the uplink ports.	
STP runs on the uplink ports.	Switch Mode
This is the default mode of operation.	

Reference:

https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-computing/whitepaper_c11-701962.html

Font Highly Voted 1 year, 4 months ago

I think :

End-Host Mode :

- This is the default mode opération
- link active-active

Switch mode :

- STP
- MAC

End-host mode features include:

- Spanning Tree Protocol is not run on both the uplink ports and the server ports.
- MAC address learning occurs only on the server ports; MAC address movement is fully supported.
- Links are active-active regardless of the number of uplink switches.
- The system is highly scalable because the control plane is not occupied.

Ethernet switching mode features include:

- Spanning Tree Protocol is run on the uplink ports per VLAN as defined by Cisco® Per-VLAN Spanning Tree Plus (PVST+)
- Configuration of Spanning Tree Protocol parameters (bridge priority, hello timers, etc.) is not supported.
- VLAN Trunk Protocol (VTP) is not supported.

- MAC address learning and aging occur on both the server and uplink ports as in a typical Layer 2 switch.
- Upstream links are blocked according to Spanning Tree Protocol rules.

upvoted 7 times

🗨️ 👤 **AboudMokh** Most Recent 5 months, 3 weeks ago

end-host mode is default

mac address learning uplink --> switch mode.

upvoted 4 times

🗨️ 👤 **Ben1009** 1 year, 5 months ago

the answer is incorrect

end-host mode is default

mac address learning uplink --> switch mode.

upvoted 4 times

🗨️ 👤 **gemini1980** 1 year, 9 months ago

from link refer.

Ethernet switching mode features include:

- Spanning Tree Protocol is run on the uplink ports per VLAN as defined by Cisco® Per-VLAN Spanning Tree Plus (PVST+)
- Configuration of Spanning Tree Protocol parameters (bridge priority, hello timers, etc.) is not supported.
- VLAN Trunk Protocol (VTP) is not supported.
- MAC address learning and aging occur on both the server and uplink ports as in a typical Layer 2 switch.
- Upstream links are blocked according to Spanning Tree Protocol rules.

i think switching mode is stp on uplink port and mac address learning on uplink

upvoted 3 times

🗨️ 👤 **gemini1980** 1 year, 9 months ago

Note: This document does not discuss the operating mode available on the Fibre Channel ports.

from link refer

For Ethernet, end-host mode is the default mode of operation. A change in the operating mode requires a fabric interconnect reboot to effect the change (Figure 2).

so the default mode is end-host mode. this answer is incorrect.

upvoted 2 times

Which two statements describe Ethernet switching mode on Cisco UCS 6100 Series Fabric Interconnects? (Choose two.)

- A. STP runs on the uplink ports of each VLAN, as defined by PVSTP+.
- B. STP runs on the uplink ports of each VLAN, as defined by MSTP.
- C. The STP parameter configuration is unsupported.
- D. The STP parameter configuration is supported.
- E. STP runs on the server ports.

Correct Answer: AC

Reference:

https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-computing/whitepaper_c11-701962.html

🗨️ 👤 **NeoTheOne** 4 months ago

AC is the correct answer

<https://community.cisco.com/t5/unified-computing-system/is-it-possible-to-config-ucs-stp-priority-in-quot-switch-mode/td-p/1747360>
upvoted 2 times

🗨️ 👤 **AboudMokh** 5 months, 3 weeks ago

Correct

upvoted 1 times

A Cisco UCS instance has four interfaces on a UCS VIC. Where on the fabric interconnect does each interface terminate?

- A. vPC
- B. virtual interface
- C. port channel
- D. physical port

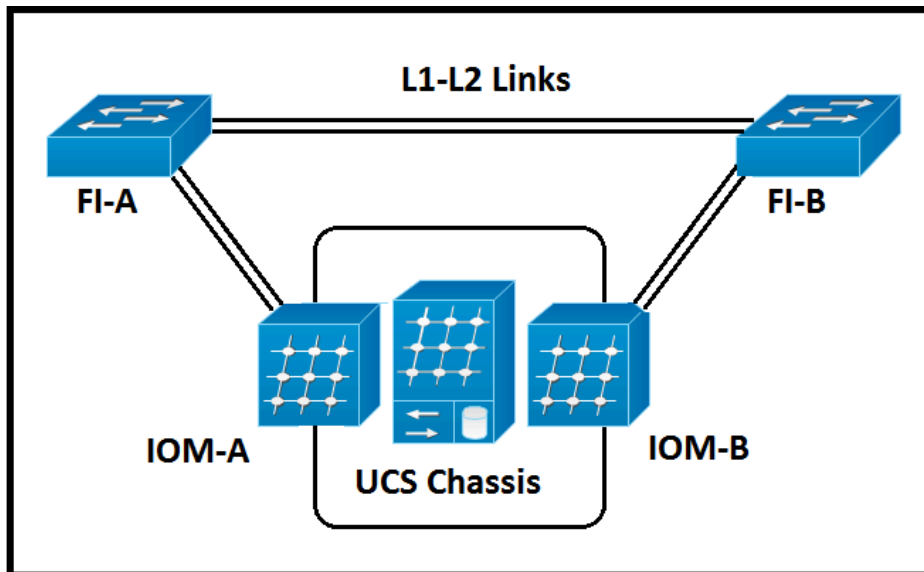
Correct Answer: *B*

Reference:

<https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-6300-series-fabric-interconnects/cisco-unified-fabric.pdf>

Currently there are no comments in this discussion, be the first to comment!

Refer to the exhibit.



A Cisco UCS chassis that has two IOMs is discovered. The chassis discovery policy creates a fabric port channel. What is the result of this policy?

- A. Cisco UCS Manager creates two separate fabric port channels.
- B. Cisco UCS Manager creates a single fabric port channel.
- C. The port channel is disabled until an administrator enables it manually.
- D. Each uplink that connects an IOM is configured as a discrete link.

Correct Answer: A

Community vote distribution

A (100%)

lurker8000 3 months, 3 weeks ago

I think the answer here is B.

upvoted 1 times

vtache 1 year, 8 months ago

Selected Answer: A

Fabric Port Channels

Fabric port channels allow you to group several of the physical links from an IOM to a fabric interconnect into one logical link for redundancy and bandwidth sharing. As long as one link in the fabric port channel remains active, the fabric port channel continues to operate.

If the correct hardware is connected, fabric port channels are created by Cisco UCS Manager in the following ways:

During chassis discovery according to the settings configured in the chassis discovery policy.

After chassis discovery according to the settings configured in the chassis connectivity policy for a specific chassis.

For each IOM there is a single fabric port channel. Each uplink connecting an IOM to a fabric interconnect can be configured as a discrete link or included in the port channel, but an uplink cannot belong to more than one fabric port channel. For example, if a chassis with two IOMs is discovered and the chassis discovery policy is configured to create fabric port channels, Cisco UCS Manager creates two separate fabric port channels: one for the uplinks connecting IOM-1 and another for the uplinks connecting IOM-2.

upvoted 1 times

DRAG DROP -

Drag and drop the configurations from the left onto the correct policies on the right.

Select and Place:


Set the number of queues.	vNIC/vHBA Placement Policy <div style="border: 1px solid black; height: 20px; width: 100%;"></div> <div style="border: 1px solid black; height: 20px; width: 100%;"></div>
Set the RSS hash value.	
Assign each vCon to a physical adapter.	Ethernet Adapter Policy <div style="border: 1px solid black; height: 20px; width: 100%;"></div> <div style="border: 1px solid black; height: 20px; width: 100%;"></div>
Assign adapter order.	

Correct Answer:

Set the number of queues.	vNIC/vHBA Placement Policy <div style="border: 1px solid black; padding: 5px; text-align: center;">Assign each vCon to a physical adapter.</div> <div style="border: 1px solid black; padding: 5px; text-align: center;">Assign adapter order.</div>
Set the RSS hash value.	
Assign each vCon to a physical adapter.	Ethernet Adapter Policy <div style="border: 1px solid black; padding: 5px; text-align: center;">Set the number of queues.</div> <div style="border: 1px solid black; padding: 5px; text-align: center;">Set the RSS hash value.</div>
Assign adapter order.	

Reference:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/gui/config/guide/2-2/b_UCSM_GUI_Configuration_Guide_2_2/configuring_network_related_policies.html#task_7F03A03C719A4A44B8ACFAD73AADC73C

 **Magneto** 6 months, 2 weeks ago

Correct

upvoted 2 times


An engineer is designing a Multichassis EtherChannel topology in which two switches must appear as a single device to a third downstream switch? Which two technologies meet this requirement? (Choose two.)

- A. HSRP
- B. VSS
- C. vPC
- D. 802.1q
- E. FEX

Correct Answer: *BC*

Reference:

https://www.cisco.com/c/dam/en/us/products/collateral/switches/nexus-7000-series-switches/C07-572835-00_NX-OS_vPC_DG.pdf

 **Magneto** 6 months, 2 weeks ago

B and C

upvoted 1 times



What is an advantage of Cisco HyperFlex Edge as compared to Cisco HyperFlex Standard Clusters?

- A. Cisco Intersight automatically creates an invisible cloud witness.
- B. Cisco UCS 5108 Blade Server chassis provides SEEPROM connectivity to solve split-brain scenarios.
- C. Cisco UCS Manager provides service profile portability across nodes.
- D. Cisco HyperFlex Edge provides native container support through the KVM hypervisor.

Correct Answer: A

Reference:

<https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/whitepaper-c11-741999.pdf>

  **Magneto** 6 months, 2 weeks ago

Correct

upvoted 1 times


Which approach splits a pair of clustered fabric interconnects into two standalone devices?

- A. Run the disable fips-mode command on each fabric interconnect.
- B. Run the connect local-management command on each fabric interconnect.
- C. Erase the configuration of both fabric interconnects.
- D. Unplug the Layer 1 and Layer 2 interfaces between the devices.

Correct Answer: D

Community vote distribution

C (100%)

 **FomaLermont85** Highly Voted 9 months, 3 weeks ago

I suppose C, because of <https://kral2.fr/crash-test-what-happen-when-you-loose-l1l2-on-an-ucsm-cluster/>
upvoted 5 times

 **FonT** Highly Voted 4 months, 1 week ago

I Confirm C : erase config

there is no way to convert back cluster into two standalone.

<https://community.cisco.com/t5/unified-computing-system/break-fabric-interconnect-cluster-into-2-standalone-fabric/td-p/4123885>

upvoted 5 times

 **k3rnelpanicpj** Most Recent 1 month, 2 weeks ago


Selected Answer: C

Only erase config can break cluster and revert FI to standalone

<https://community.cisco.com/t5/unified-computing-system-discussions/break-fabric-interconnect-cluster-into-2-standalone-fabric/td-p/4123885>

upvoted 1 times

Refer to the exhibit.



An engineer must assign multiple Fibre Channel IDs to a single port channel. Which two features must be enabled this goal? (Choose two.)

- A. SAN Pin Groups on the Cisco UCS Fabric Interconnects.
- B. NPV feature on the Cisco MDS Series.
- C. end-host mode feature on the Cisco UCS Fabric Interconnects
- D. NPIV feature on the Cisco MDS Series
- E. Smart Zoning feature on the Cisco UCS Fabric Interconnects

Correct Answer: *BD*

Community vote distribution

CD (100%)

FonT Highly Voted 3 years, 10 months ago

I confirm CD :

end-host on Fabric Interconnect to enabled NPV

NPIV enabled on SAN fabric

upvoted 10 times

FomaLermont85 Highly Voted 4 years, 3 months ago

CD should be correct answer <https://community.cisco.com/t5/unified-computing-system/npiv-npv-in-ucs/td-p/3356950>

upvoted 6 times

mojzucha Most Recent 6 months, 1 week ago

Selected Answer: CD

BD makes no sense

CD is the correct answer

upvoted 1 times

ZsF 4 years, 3 months ago

<https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2018/pdf/BRKDCN-1121.pdf>

Slide 57-58

upvoted 1 times

Which type of encoding is used on 8-Gbps links as compared to 10-Gbps links?

- A. 8-Gbps links use 64B/66B encoding, and 10-Gbps links use 8B/10B encoding.
- B. 8-Gbps links use 8B/10B encoding, and 10-Gbps links use 64B/66B encoding.
- C. 8-Gbps links and 10-Gbps links use 8B/10B encoding.
- D. 8-Gbps links and 10-Gbps links use 64B/66B encoding.

Correct Answer: B

Reference:

https://www.iol.unh.edu/sites/default/files/knowledgebase/10gec/10GbE_CI49.pdf

Community vote distribution

B (100%)

🗨️ 👤 **Wasamela** 8 months, 1 week ago

Selected Answer: B

B is correct

upvoted 1 times

🗨️ 👤 **[Removed]** 3 years, 4 months ago

, 2, 4, and 8 Gb Fibre Channel all use 8b/10b encoding. Meaning, 8 bits of data gets encoded into 10 bits of transmitted information – the two bits are used for data integrity. Well, if the link is 8Gb, how much do we actually get to use for data – given that 2 out of every 10 bits aren't "user" data? FC link speeds are somewhat of an anomaly, given that they're actually faster than the stated link speed would suggest. Original 1Gb FC is actually 1.0625Gb/s, and each generation has kept this standard and multiplied it. 8Gb FC would be 8×1.0625 , or actual bandwidth of 8.5Gb/s. $8.5 \times .80 = 6.8$. 6.8Gb of usable bandwidth on an 8Gb FC link.

10GE (and 10G FC, for that matter) uses 64b/66b encoding. For every 64 bits of data, only 2 bits are used for integrity checks. While theoretically this lowers the overall protection of the data, and increases the amount of data discarded in case of failure, that actual number of data units that are discarded due to failing serialization/deserialization is minuscule. For a 10Gb link using 64b/66b encoding, that leaves 96.96% of the bandwidth for user data, or 9.7Gb/s.

upvoted 4 times

Which Cisco HyperFlex feature provides virtual machine-level cost scaling for desktop virtualization growth?

- A. HyperFlex Edge support
- B. encryption offload cards
- C. dedicated compute nodes
- D. fabric interconnects

Correct Answer: A

Community vote distribution

C (100%)

🗨️ **FomaLermont85** Highly Voted 4 years, 3 months ago

C should be correct answer because of VDI VMs are very well deduplicated content, but VDI consumes compute resources.
upvoted 7 times

🗨️ **LemonSqueezy** 3 years, 10 months ago

Agreed, and "Edge Support" has nothing to do with VDI?
upvoted 2 times

🗨️ **FonT** Highly Voted 3 years, 10 months ago

I confirm C : compute nodes

VDI need GPU, CPU an RAM not a lot of storage. compute nodes help to reduce cost on scaling
upvoted 5 times

🗨️ **Magneto** 3 years, 6 months ago

You are correct
upvoted 1 times

🗨️ **Wasamela** Most Recent 8 months, 1 week ago

Selected Answer: C

Dedicated compute nodes, C
upvoted 1 times


Which two vSphere features must be included in the design of a Cisco HyperFlex systems stretched cluster? (Choose two.)

- A. HA
- B. VSS
- C. FT
- D. VDP
- E. DRS

Correct Answer: *AE*

Reference:

https://www.cisco.com/c/en/us/td/docs/hyperconverged_systems/HyperFlex_HX_DataPlatformSoftware/HyperFlex_Stretched_Cluster/3_5/b_HyperFlex_Systems_Stretched_Cluster_Guide_3_5.pdf

 **Magneto** 6 months, 2 weeks ago

A and E:

<https://www.cisco.com/c/en/us/support/docs/hyperconverged-infrastructure/hyperflex-hx-data-platform/214489-hyperflex-stretch-clusters-deployment-gu.html>

upvoted 1 times

DRAG DROP -


Drag and drop the components from the left onto the descriptions on the right that indicate what occurs to Fibre Channel fabric when that component fails.

Select and Place:

fabric interconnect	The Fibre Channel fabric fails for all of the connected Cisco UCS chassis.
FEX	The Fibre Channel fabric fails for one Cisco UCS chassis.
FEX link of the fabric interconnect	The Fibre Channel fabric fails for some of the servers within a Cisco UCS chassis.
converged network adapter	The Fibre Channel fabric fails for one server.

Correct Answer:

fabric interconnect	converged network adapter
FEX	FEX link of the fabric interconnect
FEX link of the fabric interconnect	fabric interconnect
converged network adapter	FEX

 **gemini1980** Highly Voted 9 months, 3 weeks ago

from link refer

<https://community.cisco.com/t5/data-center-documents/understanding-fabric-failure-and-failover-in-ucs/ta-p/3112884>

Understanding Fabric Failure

In a simple scenario of UCS system with a server with CNA card, following may happen:

- FI failure : results in fabric failure for all connected UCS chassis
- FEX failure : results in fabric failure for one UCS chassis
- FI-FEX link failure : results in fabric failure for some of the servers within a UCS chassis (depending on number of servers and uplinks)
- One CNA port failure : results in fabric failure for one server

In any of the above cases downtime can be eliminated by using redundant hardware and proper config.

upvoted 13 times

 **CHEdot** Most Recent 1 year ago

FI - fails for all of the connected Cisco UCS chassis

FEX link to FI - fails for one Cisco UCS chassis

FEX - fails for some servers within a Cisco UCS chassis.

CNA - fails for one server

upvoted 4 times

 **LionNS** 10 months, 1 week ago

It makes sense! On real exam, which answer is correct?

upvoted 1 times

 **Font** 4 months ago

Sorry CHEdot, I think the Gemini1980 answer is correct.

Fex failure will result a failure for all the chassis : half of SAN path will be broken
upvoted 2 times

What can be implemented on Cisco Nexus 5600 Series Switches for link redundancy within a Fibre Channel fabric?

- A. vPC+
- B. E-Trunk
- C. SAN port channel
- D. LACP port channel

Correct Answer: D


Community vote distribution

C (100%)

 **CHEdot** Highly Voted 4 years, 6 months ago


Should be C.

upvoted 9 times

 **CHEdot** 4 years, 6 months ago


Sorry, considering if it is asking "two" N5Ks in a pair. vPC+ may be the answer.

upvoted 1 times

 **jack425** 4 years, 5 months ago

they did not mention pair , am sticking with C is the correct answer

upvoted 5 times

 **FonT** 3 years, 10 months ago

I agree with C to

upvoted 3 times


 **mojzucha** Most Recent 6 months, 1 week ago

Selected Answer: C

there is no LACP in FC world

C is correct


upvoted 1 times

 **Wasamela** 8 months, 1 week ago

Selected Answer: C

Answer is C, FC is specifically designed for SANs.

upvoted 1 times

 **VelocityRaptor** 2 years, 2 months ago

I think is C.


https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5600/sw/san_switching/7x/b_5600_SAN_Switching_Config_7x/m_5500_configuring_sa

upvoted 2 times

 **johamandcheese** 3 years, 8 months ago

A and D have to do with ETH so they are out. C refers to a bundle of FC links similar to LACP in the ETH world.

upvoted 3 times

 **LemonSqueezy** 3 years, 10 months ago

I'm voting C. VPC+ is only useful when connecting FabricPath environments with VPC-environments, right?

upvoted 2 times

Which factor determines the choice of platform designing a Cisco HyperFlex infrastructure to support computer-aided design applications?

- A. network speed
- B. GPU support
- C. storage capacity
- D. processor count

Correct Answer: C

Community vote distribution

B (100%)

 **CHEdot** Highly Voted 4 years, 6 months ago

Answer is B. CAD relies on GPU.

Check "vGPU profiles" section on

<https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/whitepaper-c11-740243.pdf>

upvoted 10 times

 **gemini1980** 4 years, 3 months ago


yes i think B correct.

upvoted 2 times

 **FomaLermont85** 4 years, 3 months ago

HX220 and HX240 both support GPU, but HX240 has more storage

upvoted 1 times

 **FonT** 3 years, 10 months ago

you can plug more GPU card on HX240 M5 and have more choice

<https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/hx-240c-m5-specsheet.pdf>

<https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/hx-220c-m5-specsheet.pdf>

I agree with B ;-)


upvoted 1 times

 **mojzucha** Most Recent 6 months, 1 week ago

Selected Answer: B

It is B

upvoted 1 times

 **Wasamela** 8 months, 1 week ago

Selected Answer: B

GPU Support, answer is B

upvoted 1 times

What is the disadvantage of asynchronous storage replication as compared to synchronous storage replication in a disaster recovery design?

- A. reduction in application performance
- B. distance limitations
- C. specific backup requirements
- D. potential data loss

Correct Answer: *D*

Reference:

<https://www.evidian.com/products/high-availability-software-for-application-clustering/synchronous-replication-vs-asynchronous-replication/>

Currently there are no comments in this discussion, be the first to comment!

Where does Cisco UCS handle Fibre Channel traffic failover?

- A. on a Cisco UCS Fabric Interconnect ASIC in Fibre Channel switching mode
- B. on the host, by using multipathing software
- C. in the hardware on the Cisco UCS VIC 12xx adapter or later
- D. in the hardware on any Cisco UCS VIC adapter

Correct Answer: B

Reference:

https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-computing/whitepaper_c11-701962.html

Community vote distribution

B (100%)

🗨️ 👤 **Wasamela** 8 months, 1 week ago

Selected Answer: B

Answer is B, the Multipath software monitors the health of paths and automatically fails over traffic to a healthy path if a link or device fails.
upvoted 1 times

🗨️ 👤 **Virk64** 4 years, 6 months ago

Answer should be C
upvoted 1 times

🗨️ 👤 **abdo982** 4 years, 5 months ago

No Answer is B. that's the correct

Typically, host software—such as NIC teaming for Ethernet and EMC PowerPath or multipath I/O (MPIO) for Fibre Channel—provides failover across the two fabrics (Figure 4).

upvoted 10 times

🗨️ 👤 **johamandcheese** 3 years, 8 months ago

It depends. Each vHBA has a checkbox to enable failover to the other fabric. But best practice is to leave that option unchecked and have your end host do it (for example VMware Round Robin, etc). I would go B but with who knows with Cisco.

upvoted 1 times


Which mode must be enabled on a Cisco UCS Fabric Interconnect to connect directly to a storage array?

- A. NPIV mode
- B. FC end-host mode
- C. NPV mode
- D. FC switching mode

Correct Answer: D

Reference:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Storage-Mgmt/3-1/b_UCSM_GUI_Storage_Management_Guide_3_1/b_UCSM_GUI_Storage_Management_Guide_3_1_chapter_011110.html

 **ronnius** 6 months ago

(link mentioned was stuck) Answer D is right, see:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/gui/config/guide/141/UCSM_GUI_Configuration_Guide_1_4_1_chapter4.html

"Switch Mode

Switch mode is the traditional Fibre Channel switching mode. Switch mode allows the fabric interconnect to connect directly to a storage device. Enabling Fibre Channel switch mode is useful in POD models where there is no SAN (for example, a single Cisco UCS system connected directly to storage), or where a SAN exists (with an upstream MDS)."

upvoted 1 times

 **Magneto** 1 year, 6 months ago

D correct

upvoted 1 times

DRAG DROP -

Drag and drop the technologies from the left onto the correct descriptions on the right.

Select and Place:

iSCSI	allows the storage system to tell the client which paths are optimized
TOE	provides block-level access to the storage
MPIO	host software has visibility into different iSCSI ports
ALUA	offloads device CPU from encapsulation / decapsulation

Correct Answer:

iSCSI	ALUA
TOE	iSCSI
MPIO	MPIO
ALUA	TOE

  **missing_dll** 8 months ago

iSCSI > provides block-level..

TOE > offloads device CPU..

MPIO > host software has..

ALUA > allows the storage..

upvoted 3 times

Which two naming formats identify target or initiator iSCSI nodes? (Choose two.)

- A. WWPN
- B. EUI
- C. IQN
- D. WWN
- E. OUI

Correct Answer: CD

Reference:


https://www.cisco.com/en/US/docs/storage/san_switches/mds9000/sw/rel_3_x/configuration/guides/cli_3_3_1/ciscsi.html

Community vote distribution

BC (100%)

 **m_kar** Highly Voted 4 years, 5 months ago


It should be B and C
upvoted 8 times

 **LionNS** 4 years, 4 months ago

You are correct! WWN belongs to FC addressing.
upvoted 1 times

 **Wasamela** Most Recent 8 months, 1 week ago

Selected Answer: BC
B-EUI and C-IQN are the correct answer.
upvoted 1 times

 **FonT** 3 years, 10 months ago

I confirm B & C : EUI and IQN

<https://docs.vmware.com/en/VMware-vSphere/5.5/com.vmware.vsphere.storage.doc/GUID-686D92B6-A2B2-4944-8718-F1B74F6A2C53.html>

upvoted 1 times

 **gemini1980** 4 years, 3 months ago

in real exam no have choice WWN also.
upvoted 1 times


What influences the number of ISLs that are required between Cisco MDS Series switches?

- A. FCoE protocol
- B. type of storage
- C. end-to-end oversubscription ratio
- D. N Port virtualization

Correct Answer: C

Reference:

<https://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9700-series-multilayer-directors/white-paper-c11-729697.html>

 **Magneto** 6 months, 1 week ago

The number of ISLs required between Cisco MDS switches will depend on the desired end-to-end oversubscription ratio. The storage port oversubscription ratio from a single storage port to multiple servers can be used to help determine the number of ISLs needed for each edge-to-core connection.

upvoted 2 times

Which type fields can an engineer use to identify an iSCSI participant? (Choose two.)

- A. type
- B. preshared key
- C. port number
- D. identifying tag
- E. hostname

Correct Answer: *CE*


Reference:

<https://en.wikipedia.org/wiki/ISCSI>

 **lurker8000** 3 months, 3 weeks ago

Why note identifying tag?


upvoted 1 times

 **LionNS** 4 years, 4 months ago

Source is Wikipedia...

oO

upvoted 1 times

 **Magneto** 3 years, 6 months ago

Its correct

upvoted 1 times

Refer to the exhibit.

```
Switch(config-sys-qos)# show policy-map system
  Type network-qos policy-maps
=====
policy-map type network-qos system
  class type network-qos class-fcoe
    match qos-group 1
    pause no-drop
    mtu 2000
  class type network-qos class-default
    match qos-group 0
    mtu 1500
Service-policy (qos) input: system
  policy statistics status: disabled
  Class-map (qos) : class fcoe (match-any)
    Match : cos 3
    set qos-group 1
  Class-map (qos) : class-default (match-any)
    Match: any
    set qos-group 0
```

The FCoE packets fail to be forwarded through the switch. What is the minimum MTU QoS requirement for the FCoE to work?

- A. 9000 bytes
- B. 2000 bytes
- C. 2158 bytes
- D. 2240 bytes

Correct Answer: B

Community vote distribution

C (100%)


 **CHEdot** Highly Voted 4 years, 6 months ago

C is correct. This policy map shows a wrong MTU config.

Check "12. Do I need to enable Jumbo frames for FCOE?" in

<https://community.cisco.com/t5/data-center-documents/fibre-channel-over-ethernet-fcoe-questions-and-answers-from-live/ta-p/3108755>


upvoted 7 times

 **Wasamela** Most Recent 8 months, 1 week ago

Selected Answer: C

2112 bytes (Fibre Channel payload) + 4 bytes (FCB header) + 2 bytes (FCop header) = 2158 bytes Answer is C.

upvoted 1 times

 **bentoj** 4 years, 1 month ago

Answer: C

The minimum MTU is 2158 bytes

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/qos/521_n1_1/b_5k_QoS_Config_521N11/b_5k_QoS_Config_521N11_chapter.

upvoted 2 times

 **gemini1980** 4 years, 3 months ago

from link refer

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/qos/521_n1_1/b_5k_QoS_Config_521N11/b_5k_QoS_Config_521N11_chapter.

The FCoE system class (for Fibre Channel and FCoE traffic) has a default MTU of 2158 bytes. This value cannot be modified.

so i think C is correct.

upvoted 3 times

Which Ethernet to Cisco Fabric Interconnect connectivity configuration is supported when FCoE is implemented?

- A. VSS that has unified uplinks
- B. Ethernet and unified ports on the same port channel
- C. vPC that has unified uplinks
- D. cross-connections with pinning

Correct Answer: C

Reference:

<https://www.cisco.com/c/en/us/support/docs/servers-unified-computing/ucs-manager/116188-configure-fcoe-00.html#anc11>

Community vote distribution

D (100%)

🗉 👤 **Wasamela** 8 months, 1 week ago

Selected Answer: D

D is correct.

upvoted 1 times

🗉 👤 **alex010191** 3 years, 9 months ago

i think D.

FCoE Uplink with vPC supported, but no vPC that has unified uplinks

<https://www.cisco.com/c/en/us/support/docs/servers-unified-computing/ucs-manager/116188-configure-fcoe-00.html#anc7>

upvoted 4 times

🗉 👤 **ZsF** 4 years, 3 months ago

it should be D

upvoted 4 times

🗉 👤 **LemonSqueezy** 3 years, 10 months ago

Why wouldn't a VPC be supported with FCoE? I'm not sure but I like the C answer better...

upvoted 1 times

🗉 👤 **Font** 3 years, 10 months ago

I agree with D.

section "Supported Configurations" of link.

VPC could be only on ethernet no unified and the corss connection with pinning is supported

upvoted 3 times


Which two configuration methods can be used to deploy QoS to Cisco MDS Series switches? (Choose two.)

- A. zone-based QoS
- B. multiple QoS policies that match individual devices
- C. VSAN-based QoS
- D. FCoE VLAN-based QoS
- E. multiple QoS policies that match multiple devices

Correct Answer: AC

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/7_3/configuration/qos/cisco_mds9000_qos_config_guide/configuring_qos.html

 **ronnius** 6 months ago

correct link

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/6_2/configuration/guides/qos/nx-os/qos_bk/qos.html

"While you can configure both zone-based QoS and VSAN-based QoS configurations in the same switch ... "

upvoted 3 times

 **Magneto** 1 year, 6 months ago

VSAN and Zone-Based QoS

upvoted 3 times



An engineer is designing a data center that uses FCIP. How is QoS configured in this environment?

- A. A separate DSCP value and CoPP must be configured for all FCIP traffic and another for all of the FCIP in the data connection.
- B. A separate DSCP value must be configured for all FCIP VSAN traffic.
- C. A separate DSCP value and priority queuing must be configured for all FCIP traffic.
- D. Two separate DSCP values must be configured, one for all of the FCIP frames in the control TCP connection and another for all of the FCIP frames in the data connection.

Correct Answer: D

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/5_0/configuration/guides/ipsvc/nxos/ipsvc/cfcip.html

  **Magneto** 6 months, 1 week ago

Two separate DSCP values must be configured, one for control frames another one for data frames.

upvoted 4 times

Which two methods mitigate congestion in a SAN network? (Choose two.)

- A. Configure ER_RDY to allow splitting of each ISL between switches into separate virtual links.
- B. Use the port-monitor command to detect slow drain devices.
- C. Configure the port channel to enable individual buffer-to-buffer credits.
- D. Configure the port monitor to allow categorization of a specific device as slow.
- E. Configure the flow control for the FC to use R_RDY.

Correct Answer: *BD*

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/8_x/config/interfaces/cisco_mds9000_interfaces_config_guide_8x/congestion_avoidance_isolation.html

 **m_kar** 11 months ago

I would say A and D

- Extended Receiver Ready—This feature allows each ISL between supporting switches to be split into four separate virtual links, with each virtual link as buffer credits. Virtual link 0 used to carry control traffic, virtual link 1 is used to carry high-priority traffic, virtual link 2 is used to carry slow devices, and v carry normal traffic.

- Congestion Isolation—This feature allows devices to be categorized as slow by either configuration command or by port monitor.

- Port monitor portguard action for Congestion Isolation—Port monitor has a new portguard option to allow the categorization of a device as slow so that flowing to the device routed to the slow virtual link.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/8_x/config/interfaces/cisco_mds9000_interfaces_config_guide_8x/congestion
upvoted 4 times

 **Font** 4 months, 1 week ago

I agree with A & D
upvoted 1 times

 **LemonSqueezy** 4 months, 1 week ago

Also, option B doesn't actually mitigate anything it just "monitors" the problem.
upvoted 1 times

DRAG DROP -


Drag and drop the SAN components from the left onto the correct design considerations on the right.

Select and Place:

fabric logins	used to logically separate a SAN fabric
ISL	must be considered for the end-to-end oversubscription ratio
VSAN	these exchanges increase as a result of server virtualization
zone	must be considered when implementing access control

Correct Answer:


fabric logins	VSAN
ISL	ISL
VSAN	fabric logins
zone	zone

 **b44c199** 3 months, 4 weeks ago

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/8_x/config/fabric/cisco_mds9000_fabric_config_guide_8x/overview.pdf
upvoted 1 times

 **muzz40000** 1 year, 7 months ago

I dont hink so
upvoted 1 times

 **FonT** 2 years, 1 month ago

I think :

Zone : Used to logically separate a SAN Fabric

ISL : must be considered [...] Aversubscription Ratio

Fabric Login : virtualization

VSAN : considered when implementing access control

Zone separate a SAN fabric in small end-point group

VSAN could be used to separate administration access

upvoted 2 times

Which two automation requirements can be met by using Puppet? (Choose two.)

- A. Push configuration updates immediately to servers.
- B. Enable clients to pull changes to configuration files.
- C. Enable clients to pull changes without installing an agent.
- D. Provide support for writing new modules in Python.
- E. Enable one-off infrastructure changes to be performed manually.

Correct Answer: AB

Reference:

https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/whitepaper_c11-740103.pdf

Community vote distribution

AB (100%)

🗨️ 👤 **CoAsT_x** 2 months ago

Selected Answer: AB

Will also go A and B. FonT says A is wrong, but the referenced link in the Answer literally says you can provision servers and that would be a push.

upvoted 1 times

🗨️ 👤 **lurker8000** 3 months, 2 weeks ago

I will go with AB

upvoted 1 times

🗨️ 👤 **FonT** 2 years, 1 month ago

Answers should be B. E.

because :

- A is wrong : no push in puppet
- B is OK
- C is wrong, you need an agent with puppet
- D is wrong, no support from cisco and it's not python language
- E I don't know, but the best choice

upvoted 1 times

Which orchestration tool should be used for the management and analytics of servers across multiple sites?

- A. Cisco Intersight
- B. Cisco Nexus Fabric Manager
- C. Cisco UCS Manager
- D. Cisco Data Center Network Manager

Correct Answer: A

Reference:

https://www.cisco.com/c/en_in/products/servers-unified-computing/intersight/index.html

Currently there are no comments in this discussion, be the first to comment!

Which authentication software natively runs within the CLI of Cisco NX-OS devices?

- A. Python
- B. Ansible
- C. Puppet
- D. Chef

Correct Answer: B

-  **CIIMII** Highly Voted 4 years, 2 months ago
I will go for Python. It is natively supported by all nexus platforms.
upvoted 6 times
-  **Wasamela** Most Recent 8 months, 1 week ago
NX-OS doesn't have any native authentication software and the provided answers are all automation tools... I see people pointing to Puppet, but Puppet itself doesn't directly provide user authentication for logging into NX-OS devices (or any other network devices for that matter). I don't think this question is valid.
upvoted 1 times
-  **gu33mis** 3 years, 1 month ago
python is an interpreted programming language which natively runs on the NX-OS device. It's NOT an "automation" software,

I believe the correct answer is Ansible, which leverages python on the NX-OS device and doesn't require any agents to be installed.
upvoted 2 times
-  **mojzucha** 3 years, 4 months ago
"Authentication" does not make sense, of course. And as it is automation software, it must be Python. Ansible software does not run natively on the NX-OS device. NX-OS can be a client device managed by Ansible, but this does not mean it runs Ansible. The same holds for Puppet. Even if there is native Puppet agent for NX-OS, it does not mean (at least for me), it runs Puppet.
upvoted 2 times
-  **msagimlar** 3 years, 10 months ago
Puppet
<https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2015/pdf/BRKDCT-2459.pdf>
upvoted 1 times
-  **LemonSqueezy** 3 years, 10 months ago
Python is a native command in NX-OS, I will also go for Python.
upvoted 2 times
-  **CHEdot** 4 years, 6 months ago
Take from the actual exam,
this question is asking "automation" not "authentication" soft in nxos.
upvoted 4 times
-  **FomaLermont85** 4 years, 3 months ago
So, should it be Puppet? NX-OS has Puppet agent built-in
upvoted 3 times
-  **ZsF** 4 years, 3 months ago
Ansible is agentless so it runs "within" the CLI..
upvoted 2 times
-  **FonT** 3 years, 10 months ago
I agree with ansible is the question Asking for "automation software" (and it make sence).

else if it's authentication it could be python

"Python scripting of client certificate authentication is supported. If the client certificate is encrypted with a passphrase, python

successfully prompts for the passphrase. However, the passphrase cannot be passed into the script due to a current limitation with the Python requests library”

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/93x/programmability/guide/b-cisco-nexus-9000-series-nx-os-programmability-guide-93x/b-cisco-nexus-9000-series-nx-os-programmability-guide-93x_chapter_010011.html

upvoted 1 times

  **f143c37** 1 year, 9 months ago

And where did you get this info..?

upvoted 1 times

Which orchestration tool provides a graphical interface to manage LAN and SAN fabrics?

- A. Cisco Virtual Topology System
- B. Cisco Data Center Network Manager
- C. Cisco Nexus Fabric Manager
- D. Cisco CloudCenter

Correct Answer: *B*

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/DCNM_OLH/SAN_Client/fmhelp/fmc.html#34479

 **Wasamela** 8 months, 1 week ago

Correct.

upvoted 1 times

DRAG DROP -

Drag the requirements from the left onto the correct policy types on the right.

Select and Place:

Support tagging policies.	Cisco Intersight Global Policy
Scrub disks on disassociation.	
Synchronize server power state.	
Create a utility partition on a SD card.	Cisco UCS Manager Local Policy
Implement Hypervisor password policies.	

Correct Answer:

Support tagging policies.	Cisco Intersight Global Policy
Scrub disks on disassociation.	
Synchronize server power state.	
Create a utility partition on a SD card.	Cisco UCS Manager Local Policy
Implement Hypervisor password policies.	

 **LemonSqueezy** Highly Voted 3 years, 10 months ago

I'm not sure but I think the response is wrong. Create SD card policy is from Intersight, and Power Sync Policy is in UCS Manager.

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Server-Mgmt/3-2/b_Cisco_UCS_Manager_Server_Mgmt_Guide_3_2/b_Cisco_UCS_Manager_Server_Mgmt_Guide_3_2_chapter_0101.html

https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/b_Intersight_Managed_Mode_Configuration_Guide/b_Intersight_Managed_Mode_Be
upvoted 7 times

 **missionflamingo** Most Recent 1 month, 1 week ago

This is a poorly written question. IMM supports all listed policies except for power sync. UCSM does not support tagging or hypervisor password policies. My guess is as follows.

IMM:

tagging

utility
hypervisor

UCSM:

power state

scrub (IIRC this was implemented later)

upvoted 1 times

  **Wasamela** 8 months, 1 week ago

I think it should be:

Interisght:

Support Tagging Policies

Synchronize server power stare

Implement hypervisor password

UCS Manager:

Scrub disks

Create a utility partition

upvoted 1 times

Which two configurations are available by using a service profile template in Cisco UCS Manager? (Choose two.)

- A. number of vHBAs
- B. ARP entries
- C. UUID assignment
- D. VXLAN membership
- E. target server

Correct Answer: AC

Reference:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/cli/config/guide/2-2/b_UCSM_CLI_Configuration_Guide_2_2/b_UCSM_CLI_Configuration_Guide_2_2_chapter_0100000.html

Currently there are no comments in this discussion, be the first to comment!



Which two types of service profile templates does Cisco UCS support? (Choose two.)

- A. initial
- B. permanent
- C. updating
- D. temporary
- E. connected

Correct Answer: AC

Reference:

https://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/141/UCSM_GUI_Configuration_Guide_141_chapter28.html#concept_C9D27CE2AA1245EA976DAA79F5BC6808

  **Wasamela** 8 months, 1 week ago

Correct.

upvoted 1 times

Which pools are used when a vNIC template is created in Cisco UCS Manager?

- A. WWPN
- B. UUID
- C. MAC
- D. IP

Correct Answer: *C*

Reference:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Network-Mgmt/3-1/b_UCSM_Network_Mgmt_Guide_3_1/b_UCSM_Network_Mgmt_Guide_3_1_chapter_0111.html

Currently there are no comments in this discussion, be the first to comment!

Refer to the exhibit.

Create vHBA Template

Name :

Description :

Fabric ID : A B

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Select VSAN :

Template Type : Initial Template Updating Template

Max Data Field Size :

WWPN Pool :

QoS Policy :

Pin Group :

Stats Threshold Policy :

Which template is used when a vHBA template pair is created for a specific server that inherits its attributes from an existing template?

- A. Secondary
- B. Initial
- C. Updating
- D. Primary

Correct Answer: B

Community vote distribution

A (100%)

CHEdot Highly Voted 4 years, 6 months ago

How about C?

Initial Template - Created vHBA settings are not update if template changes.

Updating Template - Created vHBA settings inherits template change.

Check this video from Cisco:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/videos/4-0/creating-a-vhba-template.html

upvoted 6 times

Wasamela Most Recent 8 months, 1 week ago

Selected Answer: A

Agree with Gemini1980 based on the provided link.

upvoted 1 times

gemini1980 4 years, 3 months ago

from link refer

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Network-Mgmt/3-1/b_UCSM_Network_Mgmt_Guide_3_1/b_UCSM_Network_Mgmt_Guide_3_1_chapter_0111.html

•Secondary Template— All shared configurations are inherited from the Primary template.

•Initial Template: vNICs created from this template are not updated if the template changes.

•Updating Template: vNICs created from this template are updated if the template changes

so i think A secondary is correct.

upvoted 3 times

Magneto 3 years, 6 months ago

This makes sense
upvoted 1 times

Which address pool is used when a vHBA template is created?

- A. MAC
- B. WWPN
- C. IP
- D. UUID

Correct Answer: *B*

Reference:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/gui/config/guide/2-1/b_UCSM_GUI_Configuration_Guide_2_1/b_UCSM_GUI_Configuration_Guide_2_1_chapter_011010.html

Currently there are no comments in this discussion, be the first to comment!


What are two functions of a Network Services Orchestrator in an NFV environment? (Choose two.)

- A. It operates at the resource-facing services layer and provides overall lifecycle management at the network service level.
- B. It provides an API-based northbound interface for transparent integration with systems that operate at the resource-facing services layer.
- C. It provides a standards-based southbound interface for transparent integration with systems that operate at the customer-facing services layer.
- D. It provides a standards-based northbound interface for transparent integration with systems that operate at the customer-facing services layer.
- E. It operates at the customer-facing services layer and provides the management interface for the virtual network devices.

Correct Answer: AB

Reference:

https://www.cisco.com/c/en/us/td/docs/net_mgmt/msx/3_4/solution_overview/Cisco_VMS_Sol_Overview.pdf

 **Wasamela** 8 months, 1 week ago

Correct

upvoted 1 times



Which feature provides a sandbox to convert NX-OS CLI commands to REST-based scripts on Cisco Nexus 7700 Series Switches?

- A. NX-API
- B. Postman
- C. Bash shell
- D. Guest shell

Correct Answer: A

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/programmability/guide/b_Cisco_Nexus_9000_Series_NX-OS_Programmability_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_Programmability_Guide_7x_chapter_010010.html

  **lurker8000** 3 months, 3 weeks ago

A, NX-API

upvoted 1 times

  **damncei** 2 years, 5 months ago

<https://developer.cisco.com/docs/nx-os/#!nx-api-cli-developer-sandbox>. A

upvoted 3 times

  **mojzucha** 3 years, 4 months ago

If A says NX-API Developer Sandbox, than ok, but NX-API itself does not do that. However, there is no better option. Guestshell does not do that either.

upvoted 1 times

  **maximoh** 3 years, 5 months ago

It's D

upvoted 1 times


An engineer must use OTV for Layer 2 connectivity between data centers to support virtual machine mobility between the customer sites. To support this requirement, the engineer must ensure the existence of the same default gateway on both sites. Additionally, the operations team reports high bandwidth utilization on site A and wants to optimize the outbound traffic flows to use a local DC exit point.

Which feature must be used to meet these requirements?

- A. data group
- B. FHRP filter
- C. ARP filter
- D. control group

Correct Answer: B

Community vote distribution



🗨️ 👤 **Wasamela** 8 months, 1 week ago

Selected Answer: A

I think that based on the requirements, data groups offers the best solution for VM mobility and it optimizes outbound traffic (combined with PBR) by directing it through the local DC exit point.

upvoted 1 times

🗨️ 👤 **hinge** 1 year, 11 months ago

Selected Answer: B

https://www.cisco.com/c/dam/en/us/products/collateral/switches/nexus-7000-series-switches/guide_c07-728315.pdf

Another capability introduced by OTV is to filter First Hop Redundancy Protocol (FHRP - Hot Standby Router Protocol [HSRP], Virtual Router Redundancy Protocol [VRRP], and so on) messages across the logical overlay. This filtering is required to allow for the existence of the same default gateway in different locations and optimize the outbound traffic flows

upvoted 4 times

An engineer is designing a disaster recovery solution. In the event of a node failure, the solution should either direct all of the traffic to another node, or load-balance the traffic across the remaining nodes.

Which solution should be implemented in the design?


- A. active/active
- B. cold standby
- C. active/passive
- D. warm standby

Correct Answer: A

Community vote distribution


C (50%)

A (50%)

 **lurker8000** 3 months, 2 weeks ago

Selected Answer: A

I will go with A. "direct all of the traffic to another node, or load-balance the traffic across the remaining nodes." emphasis on loadbalance traffic across "remaining" nodes, so if you have active/active nodes, that still works
upvoted 1 times

 **Wasamela** 8 months, 1 week ago

Selected Answer: C



An Active-Passive solution involves a primary node actively handling traffic, and a secondary node in standby mode. If the primary node fails, the secondary node takes over, to either handle all traffic or load-balancing it across other nodes.
upvoted 1 times

A company runs mission-critical infrastructure in a single data center. Due to business growth, it plans to open more data centers in the coming years. A network engineer is asked to recommend a network technology for data center interconnection that allows scalability and seamless VM mobility while providing maximum security and fault isolation between data centers.

Which technology must be used to meet these requirements?

- A. TrustSec
- B. OTV
- C. CloudSec
- D. UDLD

Correct Answer: *B*

  **Wasamela** 8 months, 1 week ago

Correct

upvoted 1 times

An engineer must design a multitenant solution using Cisco 7709 switches. The solution must support separate routing instances, separate spanning tree domains, and separate firewalls for each tenant.

Which solution meets these requirements?

- A. vPC
- B. VDC
- C. CFS
- D. VRF

Correct Answer: *B*

 **Wasamela** 8 months ago

Correct

upvoted 1 times

An engineer is designing a Cisco Nexus 7000 Series disaster recovery site using the data center interconnect solution to provide gateway redundancy and needs to ensure that devices use gateways in their own data centers.

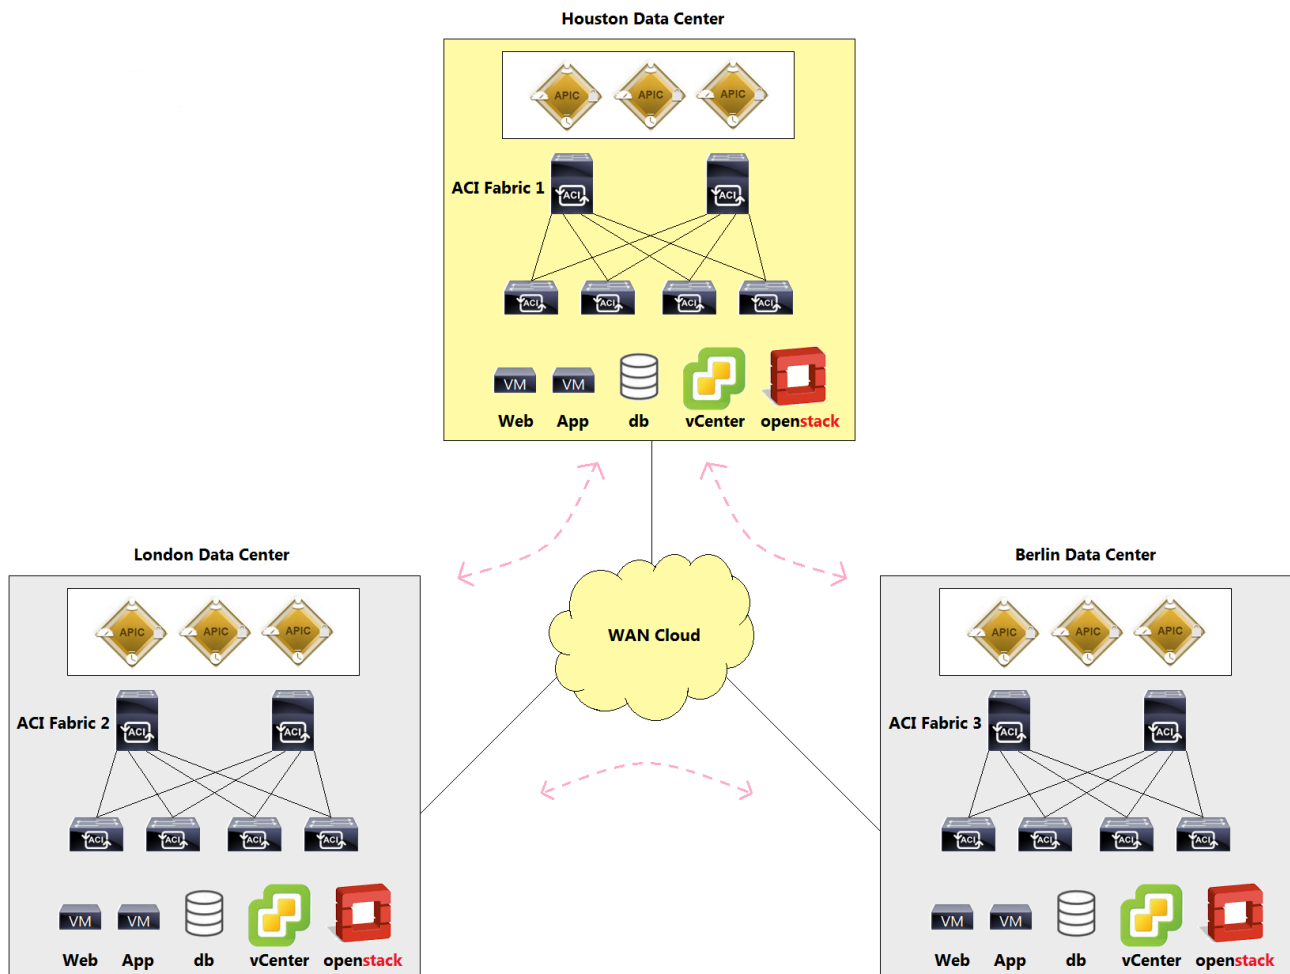
Which design approach should be used to accomplish this goal?

- A. CFS localized to its own data center
- B. vPC localized to its own data center
- C. FHRP localized to its own data center
- D. STP localized to its own data center

Correct Answer: C

Currently there are no comments in this discussion, be the first to comment!

Refer to the exhibit.



A cloud service provider deployed three geographically separated Cisco ACI data centers. One of the customers needs to host a clustered application that is spread across the three data centers. The cluster nodes that host the application need a Layer 2 topology to communicate with each other.

Which Cisco ACI topology design meets these requirements?

- A. Multi-Cloud
- B. Multi-Site
- C. Multi-Pod over OSPF EVPN
- D. Multi-Pod over MP-BGP EVPN

Correct Answer: D

Community vote distribution

B (100%)

k3rnelpanicpj 1 month, 3 weeks ago

Selected Answer: B

MultiPod can be configured on ONE ACI Fabric. In question we can find info about THREE independent Fabrics.
upvoted 1 times

lurker8000 3 months, 3 weeks ago

Agree, I will go with B as well, multi-site
upvoted 1 times

  **ciscoshare2022** 1 year, 10 months ago

wrong, Multi-Site is correct

upvoted 3 times

An engineer finishes the initial set up of a VXLAN EVPN network. The engineer is asked to plan for external connectivity that supports redundancy and extends multiple virtual routing and forwarding domains. The requirement is to use the same default gateway addressing across all leaf switches that belong to the VXLAN network.

Which two solutions must be used to meet these requirements? (Choose two.)

- A. Distributed Anycast Gateway
- B. VRF-Lite
- C. Inter-Site Network
- D. DC Interconnect
- E. Spanning Tree Protocol

Correct Answer: AD

Community vote distribution

AB (100%)

 **k3melpanicj** 1 month, 3 weeks ago

Selected Answer: AB

Anycast for same default gateway on all leaf switches,
VRF-Lite for multiple virtual routing domains
upvoted 1 times

 **lurker8000** 3 months, 3 weeks ago

I think it should be A and B.
upvoted 1 times

An engineer must configure an interface for managing and orchestrating a Cisco Nexus device. The interface must be operational if the data and control planes crash, support IPv4 and IPv6 addressing, and have access to its own VRF instance.

Which interface should be used to meet these requirements?

- A. physical interface
- B. interface VLAN
- C. management interface
- D. loopback interface

Correct Answer: C

Currently there are no comments in this discussion, be the first to comment!

An engineer is implementing management access to a Cisco Nexus platform. The customer requires a solution that provides isolation between the data and control planes, in addition to management that uses encryption and authentication. The customer will manage the device using both IPv4 and IPv6.

Which two components should be used to meet these requirements? (Choose two.)

- A. mgmt0 interface
- B. SSH protocol
- C. SNMP
- D. SVI interface
- E. Telnet protocol

Correct Answer: AB

Community vote distribution

AB (100%)

🗨️ **f143c37** 3 months, 2 weeks ago

Selected Answer: AB

Like Husam said, it's A and B definitely.

upvoted 1 times

🗨️ **husam_shabeeb** 5 months ago

Selected Answer: AB

MGMT0: Provides true out-of-band management through a dedicated interface and VRF to ensure 100 percent isolation from either control plane or data plane so idont know why snmp ? should be a&b

upvoted 2 times

An engineer must design a VXLAN EVPN network. The solution should optimize broadcast, unknown unicast, and multicast (BUM) and short BUM replication, and use PIM anycast RP.

Which solution should be included in the design?

- A. head-end replication
- B. PIM sparse mode
- C. PIM dense mode
- D. ingress replication

Correct Answer: B

Community vote distribution

B (100%)

 **k3rnelpanicj** 1 month, 3 weeks ago

Selected Answer: B

PIM sparse mode.

In dense mode there is no RP

upvoted 1 times

 **altah3r** 4 months, 3 weeks ago

Selected Answer: B

https://www.cisco.com/c/en/us/td/docs/ios/solutions_docs/ip_multicast/White_papers/anycast.html

upvoted 3 times

A network engineer designs a Cisco Nexus data center network for a client-server farm using a pair of Nexus 9000 Series Switches in a vPC setup. Each server connects with two interfaces to both Nexus switches. The Nexus switches must always initiate the negotiation with the server.

Which feature should be used in this design to ensure traffic is forwarded from the server to the data center network?

- A. PAgP mode desirable
- B. LACP mode auto
- C. LACP mode active
- D. PAgP mode auto

Correct Answer: C

Community vote distribution

C (100%)



  **k3melpanicj** 1 month, 3 weeks ago

Selected Answer: C

PAgP is Cisco proprietary protocol - no on server.

LACP active always send negotiation frames

upvoted 1 times

After experiencing traffic disruptions from the failure of a single router, a customer asks an engineer to design a solution that will prevent this from occurring in the future. While examining the customer's environment, the engineer discovers that the routers are manufactured by a variety of different vendors, and they have varying amounts of CPU and memory resources. Additionally, several of the customer's applications require the ability to fine-tune the load-balancing parameters between multiple gateway routers.

Which solution should be used to meet these requirements?

- A. VRRP
- B. FHRP
- C. HSRP
- D. GLBP

Correct Answer: A

Community vote distribution

A (67%)

D (33%)

🗨️ **k3rnelpanicj** 1 month, 3 weeks ago

Selected Answer: A

VRRP is only vendor neutral and has option for loadbalancing

upvoted 1 times

🗨️ **CoAsT_x** 2 months ago

Selected Answer: A

It has to be A. GLBP is absolutely a Cisco proprietary protocol.

upvoted 1 times

🗨️ **Wasamela** 8 months ago

Selected Answer: D

GLBP is the answer as it is vendor neutral, provides scalability for heterogeneous environments, and ability to configure load-balancing parameters effectively

upvoted 1 times

🗨️ **nmac** 8 months ago

Should be GLBP:

GLBP: Gateway Load Balancing Protocol (GLBP) allows you to use several default gateways to forward traffic upstream for server subnets, and therefore use all upstream links. However, returning traffic in the direction toward the servers usually travels only across one of these gateways. GLBP manages load distribution between default gateways and is configurable.

upvoted 1 times

An engineer must build a VXLAN flood-and-learn network. The customer requires the solution to be based on the IETF documented practice in RFC 7348. The solution must use broadcast, unknown unicast, and multicast (BUM) and include the efficient BUM and short BUM replication modes that use only shared trees.

Which technology meets these requirements?

- A. MVPN
- B. PIM sparse mode
- C. PIM SSM
- D. Bidirectional PIM

Correct Answer: D

Community vote distribution

D (100%)



 **f143c37** 9 months ago

Selected Answer: D

Sparse mode doesn't only use Shared trees, but also Shortest Path Trees. Bi-directional PIM is what uses ONLY shared trees.
upvoted 4 times

An engineer must secure the payload traversing across the DCI. The solution must support confidentiality and integrity and it must be open standard. The solution must encrypt Ethernet frames regardless of the upper layer protocol used.

Which technology must be used to support these requirements?

- A. GRE
- B. OTV
- C. MACsec
- D. IPsec

Correct Answer: C

  **ciscoshare2022** 4 months, 2 weeks ago

For

deployments requiring encryption and the capability of leveraging an Ethernet transport (public or private), MACsec offers a simplified, line-rate, per port encryption option for secure next-generation deployments.

upvoted 1 times

  **ciscoshare2022** 4 months, 2 weeks ago

<https://www.cisco.com/c/dam/en/us/td/docs/solutions/Enterprise/Security/MACsec/WP-High-Speed-WAN-Encrypt-MACsec.pdf>

upvoted 1 times

An engineer must interconnect two geographically separated data centers. The service provider offers a Layer 2 Metro Ethernet link. The strict security policy requires that the information that is shared between two sites must be sent encrypted. The application deployed between two sites requires a line-rate throughput regardless of the packet size and speed.

Which solution meets these requirements?

- A. L2TPv2
- B. IPsec
- C. MACsec
- D. L2TPv3

Correct Answer: C

Community vote distribution

C (83%)

B (17%)

🗨️ 👤 **k3rnelpanicj** 1 month, 3 weeks ago

Selected Answer: C

L2, vendor neutral and wire speed are only possible for MACSec
upvoted 1 times

🗨️ 👤 **CoAsT_x** 2 months ago

Selected Answer: C

I too think it is C. MAC sec. It is already Layer 2 by the metro link
upvoted 1 times

🗨️ 👤 **Wasamela** 8 months, 1 week ago

Selected Answer: C

MACsec can encrypt traffic over a Layer 2 Metro Ethernet link without support from the ISP. Answer is C
upvoted 3 times

🗨️ 👤 **Bob2203** 1 year, 9 months ago

Selected Answer: B

I'd go with IPsec as MACsec is a Hop to Hop encryption and therefore would need support from ISP.
upvoted 1 times

🗨️ 👤 **ciscoshare2022** 1 year, 10 months ago


For
deployments requiring encryption and the capability of leveraging an Ethernet transport (public or private), MACsec offers a simplified, line-rate, per port encryption option for secure next-generation deployments.
upvoted 1 times

An engineer configured the OSPF protocol in the vPC topology. During the catastrophic failure of one of the vPC member switches, the traffic was routed to a black hole. This route was caused by the long convergence time caused by the delay of the vPC member switch reboot. After the failure, the engineer rebooted both switches, but the secondary switch powered on before the primary one. It caused all the vPCs to be shut down due to consistency-check violations.

Which two features must be used to prevent these situations from happening in the future? (Choose two.)

- A. peer-gateway
- B. ip arp synchronize
- C. delay restore
- D. auto-recovery
- E. system priority

Correct Answer: *CD*

  **Wasamela** 8 months, 1 week ago

Correct

upvoted 1 times

A network engineer must select a high availability feature for their data center. The solution must ensure network resilience, reduce network instability for Layer 3 routing protocols, and meet these requirements:

- The device must notify its neighbors when the control plane is undergoing a restart.
- The solution must suppress routing flaps in case of stateful switchovers.

Which solution must be implemented?

- A. BFD
- B. HSRP
- C. NSF
- D. ISSU

Correct Answer: C

 **Wasamela** 8 months, 1 week ago

Correct, Non-Stop Forwarding (NSF) provides a solution for network resilience, reduced Layer 3 routing instability, control plane restart notification, and stateful switchover flap suppression.

upvoted 2 times

An engineer must logically partition two Cisco Nexus 7000 Series Switches into two environments called Test_1 and Production_1. The environments must use only F2-Series line cards and be limited to a maximum of 265 port-channels. Which Cisco Nexus feature must be used?

- A. vPC
- B. VDC
- C. vFC
- D. VRF

Correct Answer: D

Community vote distribution

B (100%)

🗨️ 👤 **Wasamela** 8 months, 1 week ago

Selected Answer: B

Virtual Device Context (VDC) is the right answer
upvoted 1 times

🗨️ 👤 **ciscoshare2022** 1 year, 10 months ago

vdc is correct
upvoted 4 times

An engineer is implementing a storage area network where an NPV is used. The engineer notices that the servers in a storage farm are experiencing heavy traffic contention. The engineer needs to migrate traffic from heavily loaded server interfaces to interfaces with less load. When the new links are deployed, the existing traffic is not rebalanced as expected. Which feature must be used to meet these requirements?

- A. NPIV
- B. multiple VSAN support
- C. DPVM
- D. disruptive load balancing

Correct Answer: A

Community vote distribution

D (100%)

🗨️ 👤 **ciscoshare2022** Highly Voted 👍 1 year, 10 months ago

Disruptive Auto Load Balancing of Server Logins across NP Links

FC NPV supports disruptive load balancing of server logins. When disruptive load balancing is enabled, FC NPV redistributes the server interfaces across all available NP uplinks when a new NP uplink becomes operational. To move a server interface from one NP uplink to another NP uplink, FC NPV forces reinitialization of the server interface so that the server performs a new login to the core switch.

FC NPV supports disruptive load balancing of server logins. When disruptive load balancing is enabled, FC NPV redistributes the server interfaces across all available NP uplinks when a new NP uplink becomes operational. To move a server interface from one NP uplink to another NP uplink, FC NPV forces reinitialization of the server interface so that the server performs a new login to the core switch.

Only server interfaces that are moved to a different uplink are reinitialized. A system message is generated for each server interface that is moved.

upvoted 5 times

🗨️ 👤 **CoAsT_x** Most Recent 🕒 2 months ago

Selected Answer: D

Ya, per the comments, and own expertise, would agree that disruptive load balancing feature would be required.

upvoted 1 times

🗨️ 👤 **lurker8000** 3 months, 3 weeks ago

D. Disruptive LB

upvoted 2 times

A customer needs a design for a new data center network that meets these requirements:

- Logically partition the environment into separate security domains.
- Cut costs related to purchasing additional equipment.
- Separate routing processes, TCAM, QoS tables, and data planes.

Which technology must be used?

- A. VLAN
- B. vPC
- C. VDC
- D. VRF

Correct Answer: D

Community vote distribution

C (100%)

🗨️ 👤 **Wasamela** 8 months, 1 week ago

Selected Answer: C

Virtual Device Context (VDC) sounds like the closes solution.
upvoted 2 times

🗨️ 👤 **Fcpoultry** 1 year, 4 months ago

The request is which technology must be used, however it doesn't mean that it will resolve all requirements, so with this clarification lets see the options:

- A. VLAN - it will not cut costs related to purchasing additional equipment, also it will not separate routing processes...
- B. vPC - it will not meet any requirement;
- C. VDC - the only requirement that it will not meet is the TCAM, however this technology must be used to attend most of requirements. So, C is the best option.
- D. VRF - it will not meet: TCAM, QoS table, and data planes

upvoted 2 times

🗨️ 👤 **f143c37** 1 year, 9 months ago

TCAM is shared between all VDCs. Thus if we follow this requirement, only vPC is correct. But that goes against the second requirement.
upvoted 1 times

🗨️ 👤 **ciscoshare2022** 1 year, 10 months ago

for item 3 "Separate routing processes, TCAM, QoS tables, and data planes."
only vdc is correct
upvoted 1 times

🗨️ 👤 **f143c37** 1 year, 9 months ago

TCAM is shared between VDCs, so no, vdc isn't the correct option.
upvoted 1 times

A company has several data centers around the world that must be interconnected over Layer2. The service provider supports MPLS services to all the locations. The proposed solution must be vendor agnostic and scalable enough to support the rapid growth of the company. Which data center technology is recommended to achieve this goal?

- A. OTV
- B. VXLAN EVPN
- C. FabricPath
- D. EoMPLS

Correct Answer: B

Community vote distribution

B (100%)

 **k3rnelpanicpj** 1 month, 3 weeks ago

Selected Answer: B

Two possible:

VXLAN EVPN and EoMPLS

Question is about DATA CENTER technology - in that case VXLAN EVPN

EoMPLS is provider solution

upvoted 1 times

A cloud service provider provisioned four active/active data centers in various locations within the country at 50 miles between each data center. The data centers must provide redundant network infrastructure and always be available to customers.

Which two data center design steps must be used to meet these requirements? (Choose two.)

- A. Provide application hosting services to customers from the closest data center.
- B. Implement active cluster in data centers 1 and 2 and backup in 3 and 4.
- C. Use asynchronous replication between data centers 1 and 2.
- D. Deploy database services in data center 1 and application layer to other data centers.
- E. Configure application data replication to all backup data centers.

Correct Answer: *AE*

Currently there are no comments in this discussion, be the first to comment!

An engineer must design a DNS service available to multiple network zones as a shared service. The network zones are deployed as VRFs within the data center network and no firewall is available for communication between VRFs.

Which protocol is needed to implement the shared services?

- A. BGP
- B. OSPF
- C. ISIS
- D. VXLAN

Correct Answer: D

Community vote distribution

B (50%) D (50%)

🗨️ **lurker8000** 3 months, 2 weeks ago

Selected Answer: D

VXLAN is correct, there was no VRF isolation requirement. The requirement was a shared service that had DNS that should be reachable to everything in the network.

upvoted 1 times

🗨️ **Fcpoultry** 1 year, 1 month ago

VXLAN is not an appropriate solution, as you'll allow the traffic between all VRF as there is no control access via VXLAN, so there is no restriction among all VRFs as one routing table.

upvoted 2 times

🗨️ **Fcpoultry** 1 year, 1 month ago

Selected Answer: B

Route replication is supported for static, EIGRP, and OSPF routes. It is not possible to replicate routes to and from BGP, but that is not an issue because the BGP import and export method of copying routes between VRFs is available in a virtual network.

<https://www.cisco.com/en/US/docs/ios-xml/ios/evn/configuration/xe-3sg/evn-shared-svcs.html>

upvoted 1 times

🗨️ **b44c199** 3 months, 1 week ago

As you mentioned, it refers to BGP

upvoted 1 times

An engineer must design a secondary site to prevent a catastrophic primary site failure. The applications must communicate across sites via the data link layer of the OSI model. The Layer 3 transport that is available for this solution consists of 10-Gb connections over Layer 3.

Which two solutions meet these requirements? (Choose two.)

- A. OTV
- B. vPC
- C. OSPF
- D. IS-IS
- E. VXLAN EVPN

Correct Answer: *AE*

Currently there are no comments in this discussion, be the first to comment!

An engineer enables the HSRP feature on a Cisco Nexus 7000 Series Switch. It must be ensured that the switch manages packets that are sent to the local vPC MAC address, remote vPC MAC address, and HSRP virtual MAC address.


Which command accomplishes this goal?

- A. peer-switch
- B. peer-gateway
- C. map-server
- D. hsrp preempt

Correct Answer: D

Community vote distribution

B (100%)

 **Bob2203** 3 months ago

Selected Answer: B

Peer gateway

upvoted 2 times

 **ciscoshare2022** 4 months, 2 weeks ago

vPC Peer Gateway and HSRP

Some third-party devices can ignore the HSRP virtual MAC address and instead use the source MAC address of an HSRP router. In a vPC environment, the packets using this source MAC address may be sent across the vPC peer link, causing a potential dropped packet. Configure the vPC peer gateway to enable the HSRP routers to directly handle packets sent to the local vPC peer MAC address and the remote vPC peer MAC address, as well as the HSRP virtual MAC address. See the Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide, Release 5.x, for more information on the vPC peer gateway.

upvoted 1 times

A network engineer must design a data center network with four Cisco Nexus 7706 switches. The switches will be installed in primary and secondary data centers in vPC domains. The two data centers must be connected via Layer 2 links and provide the maximum throughput, redundancy, and avoid port-blocking spanning-tree design.

Which deployment model meets these requirements?

- A. dual-sided single vPC
- B. single-sided dual vPC
- C. dual-sided dual vPC
- D. single-sided single vPC

Correct Answer: C

  **ciscoshare2022** 4 months, 2 weeks ago

Required Recommendations:

- Use different vPC domain-id for each vPC domain (DC1: vPC domain for aggregation, vPC for DCI. DC2: vPC domain for aggregation, vPC for DCI)
 - For each data center, interconnect the aggregation vPC domain to the DCI vPC domain using a vPC (double-sided topology)
- upvoted 1 times

An engineer is implementing a Cisco UCS environment. The requirements for the solution are for the MAC addresses to be learned only on the server-facing ports. The use of the Spanning Tree Protocol must be avoided. Also, the engineer must have deterministic traffic handling from blades to uplink ports.

Which two settings must be selected to meet these requirements? (Choose two.)

- A. end-host mode
- B. dynamic pinning
- C. NPIV mode
- D. static pinning
- E. switching mode

Correct Answer: AD

  **Wasamela** 8 months ago

Correct

upvoted 1 times



An engineer is installing a Cisco Nexus 7000 Series Switch with F2-Series I/O Module. The company's strict security policy requires that the port be in-band. The engineer must also keep these considerations in mind:

- The customer's security requirements mandate the use of a secure management mechanism that supports authentication and encryption.
- The port should support management over IP.

Which port type must be used to meet these requirements?

- A. management port
- B. console port
- C. connectivity management port
- D. Ethernet port

Correct Answer: A

  **b44c199** 3 months, 4 weeks ago

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/best_practices/cli_mgmt_guide/cli_mgmt_bp/connect.html
upvoted 1 times

An engineer is building a data center network with VXLAN EVPN. The requirement is to use multicast for the broadcast, unknown unicast, and multicast replication.

Which two VNI to multicast group mapping methods must be used to meet these requirements? (Choose two.)

- A. many NVE to one multicast group
- B. one VNI to one multicast group
- C. many VNI to one multicast group
- D. one NVE to many multicast groups
- E. many VNI to many multicast groups

Correct Answer: AC

Community vote distribution

BC (100%)

🗨️ **k3rnelpanicj** 1 month, 2 weeks ago

Selected Answer: BC

NVE is network virtual interface for VXLAN encapsulation/decapsulation - exist only one.

VNI can be mapped to multicast group in manner:

1 VNI :1 MCastGroup

Many VNI: 1McastGroup

upvoted 1 times

🗨️ **lurker8000** 3 months, 2 weeks ago

Selected Answer: BC

going with this

upvoted 1 times

🗨️ **lurker8000** 3 months, 2 weeks ago

I'm going with B and C, the question says "Which two VNI mcast group mapping"...

upvoted 1 times

🗨️ **Wasamela** 8 months, 1 week ago

Correct

upvoted 1 times

A data center uses Cisco Nexus 7000 Series Switch and contains two tenants, tenant A and tenant B.

- Tenant A uses IP addresses in the 10.2.30.41/24 to 10.2.30.64/24 range.
- Tenant B uses IP addresses in the 10.2.30.120/24 to 10.2.30.192/24 range.

The engineer must implement a solution that:

- Isolates the traffic for each tenant as separate routing data and control planes.
- Physically partitions the Cisco Nexus device into separate multiple logical devices with separate management, failure domain isolation, and resource management.

Which two technologies meet these requirements? (Choose two.)

- A. VDC
- B. vPC
- C. VXLAN
- D. VLAN
- E. VRF

Correct Answer: *AE*

Currently there are no comments in this discussion, be the first to comment!

An engineer must design a solution based on Cisco Nexus 7709 Switches. The solution must connect test and development network segments and separate the enterprise applications from the production network. The network security team requires routing instances and the IP stack to be completely separated from each other and for traffic between the environments to undergo a deep packet inspection. The development team also requests the ability to allocate additional physical interfaces from the Cisco Nexus 7709 device in each environment.

Which solution meets these requirements?

- A. 1. Create a VRF for the testing and development environments.
2. Use ACLs to control traffic between the two environments.
- B. 1. Enable new routing instances for the testing and development environments in the same VDC.
2. Control reachability by using firewall access policies.
- C. 1. Assign a separate VDC for the testing and development environments.
2. Place a firewall between the VDCs.
- D. 1. Create a separate VRF for the testing and development environments.
2. Control reachability by using firewall access policies.

Correct Answer: C

Currently there are no comments in this discussion, be the first to comment!

A network architect proposes a distinct Fibre Channel fabric to be used for a Cisco UCS blade server that hosts critical applications.


Which action should be implemented for the Fibre Channel traffic from the vHBA of this server to pass through the I/O Module in the Cisco UCS chassis to a specified uplink Fibre Channel port?

- A. Update the applied global QoS policy.
- B. Update the applied Fibre Channel adapter policy.
- C. Enable Persistent Binding in the vHBA policy.
- D. Include a SAN pin group in the vHBA policy.

Correct Answer: D

Community vote distribution

D (100%)

 **lurker8000** 3 months, 2 weeks ago

Selected Answer: D

agree with D.

upvoted 1 times

 **Fcpoultry** 1 year, 1 month ago

Selected Answer: D

"Cisco UCS uses SAN pin groups to pin Fibre Channel traffic from a vHBA on a server to an uplink Fibre Channel port on the fabric interconnect. You can use this pinning to manage the distribution of traffic from the servers."

https://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/141/UCSM_GUI_Configuration_Guide_141_chapter21.html

upvoted 2 times

An engineer is experiencing performance issues on a Cisco UCS blade server. The B-Series blade server contains four CPUs, most of which are idle. The engineer notices that the CPU is suffering from too many requests sent by the NIC. Additionally, the number of queues appears to be insufficient and only a single CPU is processing the network traffic.

Which policy must be used to alleviate these issues?

- A. LAN connectivity
- B. dynamic vNIC connection
- C. Ethernet adapter
- D. vNIC placement

Correct Answer: C

Community vote distribution

C (83%)

D (17%)

🗨️ 👤 **CoAsT_x** 2 months ago

Selected Answer: C

It is 1000% Ethernet adapter policies. C. non-negotiable. <https://www.cisco.com/c/en/us/products/collateral/interfaces-modules/unified-computing-system-adapters/victuning-wp.html#Transmitandreceivequeues>

upvoted 1 times

🗨️ 👤 **lurker8000** 3 months, 2 weeks ago

Selected Answer: D

D vNIC placement should be the answer here...

upvoted 1 times

🗨️ 👤 **Fcpoultry** 1 year, 1 month ago

Selected Answer: C

"RSS distributes network receive processing across multiple CPUs in multiprocessor systems. This can be:
disabled—Network receive processing is always handled by a single processor even if additional processors are available.
enabled—Network receive processing is shared across processors whenever possible."

https://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/141/UCSM_GUI_Configuration_Guide_141_chapter19.html#d4900e910_n

upvoted 4 times

A network engineer is designing a Cisco Hyperflex infrastructure. The infrastructure will be connected to an external SAN fiber channel array and must support Microsoft clustering.

Which storage connectivity method must be used to meet these requirements?

- A. Fibre Channel over Ethernet
- B. Internet Small Computer Systems Interface
- C. raw device mapping
- D. server message block

Correct Answer: C

Community vote distribution

C (75%)

A (25%)

🗨️ 👤 **CoAsT_x** 2 months, 3 weeks ago

Selected Answer: C

Answer is C

Cisco Press - Hyperconverged Infrastructure Data Centers: Demystifying HCI.
Chapter 9. Cisco HyperFlex

HyperFlex Integration with External Storage

Examples of some applications where HX integrates with legacy infrastructure include

Boot and run VMs that are stored on the legacy SAN.

Use the SAN for storage backup.

Migrate VMs and storage to the newer hyperconverged infrastructure.

Use raw device mapping (RDM) from the fibre channel array for Microsoft clustering.

upvoted 1 times

🗨️ 👤 **lurker8000** 3 months, 2 weeks ago

Selected Answer: A

A is correct

upvoted 1 times

🗨️ 👤 **lurker8000** 4 months, 1 week ago

I think this should be A.

Yes, Fibre Channel over Ethernet (FCoE) can support Microsoft Clustering when using external storage solutions. In a clustered environment, shared storage is essential for failover and redundancy, and FCoE allows you to connect to Fibre Channel storage arrays over an Ethernet network.

upvoted 1 times

🗨️ 👤 **Fcpoultry** 1 year, 1 month ago

Selected Answer: C

The appropriate option is RDM(Raw Device Mapping) for Microsoft external storage.

upvoted 2 times

An engineer installs a B-Series server in a Cisco UCS chassis that contains a pair of UCS-IOM-2208XP I/O modules. The full bandwidth of the I/O module must be utilized.

Which two elements should be used in this environment? (Choose two.)

- A. four twinax cables from each I/O module to each fabric interconnect
- B. Cisco UCS VIC 1225
- C. eight twinax cables from each I/O module to each fabric interconnect
- D. Cisco UCS M81KR
- E. Cisco UCS VIC 1280

Correct Answer: *CE*

Currently there are no comments in this discussion, be the first to comment!

A network engineer is implementing Cisco UCS Manager with Cisco UCS Central integration. The company's strict security policies require that the logical profile of virtual network adapters and logical addressing assigned to Cisco UCS servers be managed by the Cisco UCS Manager. The users must also be prevented from being able to acknowledge pending activities in the Cisco UCS Manager.

Which two profiles must be used to meet these requirements? (Choose two.)

- A. global maintenance
- B. local service
- C. local power
- D. restricted chassis
- E. public network control

Correct Answer: BE

Community vote distribution

AB (67%)

BD (33%)

🗨️ **b44c199** 3 months, 1 week ago

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-central/GUI-User-Guides/Operations/b_UCSC_Ops_Guide/b_UCSC_Ops_Guide_chapter_0111.html
upvoted 1 times

🗨️ **lurker8000** 3 months, 2 weeks ago

Selected Answer: BD

Restricted Chassis Profile: This profile limits user access to the chassis level. It prevents users from acknowledging pending activities, ensuring that only authorized personnel can manage critical operations.

Public Network Control Profile: This profile allows centralized management of virtual network adapters and logical addressing through Cisco UCS Manager. By using this profile, the company can enforce consistent network configurations and security policies across all UCS servers.

upvoted 1 times

🗨️ **lurker8000** 3 months, 3 weeks ago

I will go with Local service and restricted chassis B and D

upvoted 1 times

🗨️ **Fcpoultry** 1 year, 4 months ago

Selected Answer: AB

The A & B are the correct answers.

upvoted 2 times

A network engineer must design a Cisco HyperFlex solution based on these requirements:

- two clusters in the main data center consisting of five HyperFlex nodes
- one edge node for the remote branch
- cluster nodes that use one rack unit of space


Which two devices should be used in the remote branch cluster for this design? (Choose two.)

- A. UCS B200
- B. Gigabit Ethernet Switch
- C. Fabric Interconnect 6300
- D. HX240c
- E. HX220c

Correct Answer: BE

Community vote distribution

BE (100%)

 **f143c37** 9 months ago

Selected Answer: BE

The HX240C is 2RU and not 1RU. Below is the link for Hx220C which as stated in the document, requires a gigabit ethernet switch!

<https://www.cisco.com/c/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/hx220c-edge-m5-ds.html>

upvoted 3 times

A customer needs to deploy a three-tier network architecture that will consist of web, app, and other business-critical applications. The solution will include Cisco UCS 6324 Fabric Interconnects managing B and C series servers. The fabric interconnects will be connected to the Cisco Nexus 7706 Series Switch for network access. A storage array will be directly connected to the UCS fabric interconnects.

Which UCS deployment type must be used to meet these requirements?

- A. UCS interconnect in Ethernet switching mode
UCS interconnect in Fibre Channel switching mode
Fabric Interconnects uplinked to the Nexus device using an LACP port-channel in passive mode
- B. UCS interconnect in Ethernet end-host mode
UCS interconnect in Fibre Channel end-host mode
Fabric Interconnects uplinked to the Nexus device using an LACP port-channel in passive mode
- C. UCS interconnect in Ethernet end-host mode
UCS interconnect in Fibre Channel switching mode
Fabric Interconnects uplinked to the Nexus device using an LACP port-channel in active mode
- D. UCS interconnect in Ethernet switching mode
UCS interconnect in Fibre Channel end-host mode
Fabric Interconnects uplinked to the Nexus device using an LACP port-channel in active mode

Correct Answer: D

Community vote distribution

C (100%)

🗨️ 👤 **ham122** 9 months ago

C is correct

upvoted 1 times

🗨️ 👤 **f143c37** 9 months ago

Selected Answer: C

We're connecting a storage array directly to FIs. Which means FI (ports) have to be in FC switching mode.

upvoted 2 times

A client in a commercial building installs a Cisco blade server in the multitenant Cisco UCS blade infrastructure. To cut operating costs, the client's data center is shared with other organizations. The client asks an engineer to prevent third parties from accessing their equipment and from accessing BIOS-related information when booting a server.

Which two policies must be used to meet these requirements? (Choose two.)

- A. post error pause
- B. CDN control
- C. ASPM support
- D. front panel lockout
- E. quiet boot

Correct Answer: *DE*

Currently there are no comments in this discussion, be the first to comment!

An engineer must connect an existing Cisco UCS Fabric Interconnect to an IP-based storage array. The connection must be a port channel made directly between the fabric interconnect expansion module and the storage array. The engineer must also consider that the environment is running FCoE for storage and uses QoS to prioritize traffic.

Which port type should be used to meet these requirements?

- A. network
- B. uplink
- C. server
- D. appliance

Correct Answer: D

Community vote distribution

D (75%)

B (25%)

🗨️ **lurker8000** 3 months, 2 weeks ago

Selected Answer: B

I will go with uplink on this one since it's an IP-based storage.

upvoted 1 times

🗨️ **Fcpoultry** 1 year, 1 month ago

Selected Answer: D

It is a directly connection to the Storage array, so the port should be an appliance.

upvoted 3 times

An engineer designs a Cisco UCS solution that must provide guaranteed and deterministic bandwidth to a specific server in the environment. The solution must apply to network and storage traffic of C-Series and B-Series servers.

Which solution should be included to meet these requirements?

- A. Pin the vNICs and the vHBAs of the service profile to dedicated uplink ports.
- B. Pin the IOM ports of the service profile to dedicated server ports.
- C. Pin the server ports of the service profile to dedicated uplink ports.
- D. Pin the vNICs and vHBAs of the service profile to dedicated server ports.

Correct Answer: A

Community vote distribution

A (80%)

C (20%)

🗨️ **lurker8000** 3 months, 2 weeks ago

Selected Answer: A

yes A is correct

upvoted 1 times

🗨️ **Wasamela** 8 months ago

Selected Answer: A

Answer A should be the correct one as pinning the vNICs and vHBAs of the service profile to dedicated uplink ports work with Cisco UCS environment with C-Series and B-Series servers.

upvoted 2 times

🗨️ **Fcpoultry** 1 year, 1 month ago

Selected Answer: A

you can't pin the server ports, however you can pin the vNICs and the vHBAs to the dedicated uplinks ports.

upvoted 1 times

🗨️ **Rick099** 1 year, 4 months ago

Selected Answer: C

C is correct

upvoted 1 times

A network consultant evaluates the Cisco UCS Fabric Interconnect SAN configuration, which must meet these requirements:

- All traffic must pass over the upstream device for switching and domain IDs must be preserved in the fabric.
- Each server should be pinned to one designated upstream interface and for the traffic to pass through the pinned port.

Which Fabric Interconnect SAN operational mode meets these requirements?

- A. end host
- B. switch
- C. NPIV
- D. Fibre Channel

Correct Answer: A

Community vote distribution

A (100%)

🗨️ **lurker8000** 4 months ago

I think it should be B Switch mode:

A. End Host Mode: In End Host mode, the Cisco UCS Fabric Interconnect acts as a gateway to the SAN, but it doesn't preserve domain IDs in the fabric. This mode is typically used to simplify fabric management and reduce traffic overhead by connecting only directly to the upstream device, not allowing traditional Fibre Channel domain management.

upvoted 1 times

🗨️ **Rick099** 1 year, 4 months ago

Selected Answer: A

End-host mode is synonymous with NPV mode. This is the default Fibre Channel Switching mode.

"https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/gui/config/guide/141/UCSM_GUI_Configuration_Guide_1_4_1_chapter4.html#concept

upvoted 2 times

🗨️ **f143c37** 1 year, 9 months ago

Selected Answer: A

There's no pinning in FC switch mode.

upvoted 4 times

An engineer is implementing a Cisco UCS environment. The system must aggregate all vNICs into a single policy and consistently enforce node disassociation across the cluster. Which two UCS policies meet these requirements? (Choose two.)

- A. scrub policy
- B. vNIC policy
- C. boot policy
- D. firmware policy
- E. LAN connectivity policy

Correct Answer: AE

Community vote distribution

AE (75%)

BE (25%)

🗨️ 👤 **CoAsT_x** 2 months ago

Selected Answer: AE

Definitely LAN connectivity policy. but this enforce node disassociation across the cluster makes no sense. There is no such thing as a "vNIC policy" per se. Scrub policy is what happens to the local disks upon disassociation, wipes the disks or not. it doesn't enforce the disassociation. Server Pools do, per se, but that isn't an option here. The only thing even RELATED to disassociation is scrub policy, so....idk. this is dumb question, but i would also have to say A and E...i guess. If they are talking about a vNIC "policy" (creating a vNIC) there are no options to enforce disassociation of any sorts. DUMB. Just wondering if there is something missing in the wording with this question of this will be the question verbatim.

upvoted 1 times

🗨️ 👤 **lurker8000** 3 months, 2 weeks ago

Selected Answer: BE

I will go with B & E on this one

upvoted 1 times

🗨️ 👤 **lurker8000** 3 months, 3 weeks ago

I'll go with B and E

upvoted 1 times

🗨️ 👤 **Rick099** 1 year, 4 months ago

Selected Answer: AE

- vNIC Template This policy defines how a vNIC on a server connects to the LAN. This policy is also referred to as a vNIC LAN connectivity policy.
- Scrub Policy: This policy determines what happens to local data and to the BIOS settings on a server during the discovery process and when the server is disassociated from a service profile.

upvoted 2 times

A cloud service provider is managing day-2 operations of over 200 blade servers installed in the Cisco UCS environment. This pool of network resources is providing virtual networks to their customers. To successfully operate this environment, frequent changes to vNICs and other operational policies are required across multiple UCS domains. The operations team must streamline these repetitive tasks. Which orchestration tool must be used to meet this requirement?

- A. Cisco ACI Multisite Orchestrator
- B. Cisco DNA Center
- C. Cisco UCS Director
- D. Cisco UCS Manager

Correct Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **Sephzer** 8 months, 1 week ago

Selected Answer: C

Agreed, the answer is C

upvoted 1 times

🗨️ 👤 **f143c37** 1 year, 9 months ago

Selected Answer: C

What does ACI even have to do with this?

upvoted 4 times

An engineer deploys VMware ESXi hypervisor on blade servers. The virtual networking is managed by Cisco Nexus 1000V Series Switches.

Which configuration aligns with Cisco recommended practices for providing high availability and a maximum throughput to the VMware ESXi hosts?

- A. dual vNICs with static EtherChannel link aggregation
- B. dual vNICs with hardware fabric failover on the fabric interconnects
- C. dual vNICs with MAC pinning link aggregation
- D. dual vNICs with LACP link aggregation

Correct Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

An engineer is operating data center environment that hosts data-intensive financial applications. The applications are mostly processing HTTP/HTTPS data with large data segments, which results in a CPU contention due to the significant network processing.

Which set of Ethernet adapter policies must be selected to resolve the issue?

A. Receive Checksum Offload field set to Disabled
TCP Segmentation Offload field set to Disabled

Transmit Queues field: 32 -
Ring Size field: 256

B. Receive Checksum Offload field set to Enabled
TCP Segmentation Offload field set to Disabled

Transmit Queues field: 64 -
Ring Size field: 128

C. Receive Checksum Offload field set to Enabled
TCP Segmentation Offload field set to Enabled

Transmit Queues field: 128 -
Ring Size field: 64

D. Receive Checksum Offload field set to Disabled
TCP Segmentation Offload field set to Enabled

Transmit Queues field: 256 -
Ring Size field: 128

Correct Answer: B


Community vote distribution

C (100%)

 **lurker8000** 4 months ago

Agree to C

upvoted 1 times

 **Rick099** 1 year, 4 months ago

Selected Answer: C

- Receive Checksum Offload field: Enabled ==> The CPU sends all packet checksums to the hardware for validation. This option may reduce CPU overhead.
- TCP Segmentation Offload field = Enabled ==> The CPU sends large TCP packets to the hardware to be segmented. This option may reduce CPU overhead and increase throughput rate.

https://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/141/UCSM_GUI_Configuration_Guide_141_chapter19.html#d4900e910_n4
upvoted 1 times

An engineer must deploy a Cisco HyperFlex system at two data centers. For redundancy reasons, the setup must achieve data protection in case of an entire site failure. The data must be available on both sites at all times to keep the recovery time objective to a minimum.

Which Cisco HyperFlex deployment model meets these requirements?

- A. ROBO cluster
- B. stretched cluster
- C. edge with four nodes or more
- D. standard with eight nodes

Correct Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

A network consultant proposes that the configuration of the Cisco UCS servers be based on service templates to facilitate the UCS infrastructure operation. The service profiles must remain connected to the template and be modified when the template settings are changed.

Which configuration meets these requirements?

- A. service profile with type set to Global
- B. service profile with type set to Local
- C. service profile template with type set to Updating
- D. service profile template with type set to Initial

Correct Answer: C

Currently there are no comments in this discussion, be the first to comment!

An engineer must integrate Layer 4 to Layer 7 services within Cisco ACI Fabric. The main requirement is to have a single point of configuration for network elements and firewall policies.

Which action meets these requirements?

- A. Define the service graph model as network policy mode.
- B. Deploy the firewalls in routed mode with PBR.
- C. Deploy the firewalls in transparent mode.
- D. Define the service graph model as service policy mode.

Correct Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

A cloud provider deploys an infrastructure based on Cisco ACI Fabric, Cisco UCS, and Cisco HyperFlex Data Platform. The storage area network will be based on Cisco MDS 9000 Series Switches with Hyper-V, VMware, and KVM used as the virtualization platform. The requirement is to deploy the overall solution by using a robust automation and orchestration tool.

Which tool meets these requirements?

- A. Cisco Prime Infrastructure
- B. Cisco Network Assurance
- C. Cisco UCS Director
- D. Cisco Cloud Center

Correct Answer: C

Community vote distribution

D (100%)



 **lurker8000** 4 months ago

I think this should be C, answer is correct:

Cisco CloudCenter (now known as Cisco Intersight in its newer iterations) is designed for multi-cloud management and application lifecycle management across hybrid and multi-cloud environments. While it can work with a variety of platforms, it is more focused on application orchestration and deployment rather than the comprehensive data center automation and orchestration that UCS Director provides for infrastructure components like UCS, HyperFlex, and MDS.

upvoted 1 times

 **Wasamela** 8 months ago

Selected Answer: D

Answer should be D, as Cloud Center provides automating, deployment orchestrating and management of ACI, UCS, HyperFlex, MDS switches, and various hypervisors.

upvoted 1 times

An engineer must design a Cisco HyperFlex solution to support a virtualized environment in a single data center location. The design must consider these customer requirements:

- a large-scale cluster with eight or more nodes
- distributed high-performance file system for virtualized servers
- high fault tolerance to multiple node failures

Which action meets these objectives?

- A. Configure data destaging and deduplication.
- B. Implement a replication factor of four.
- C. Enable logical availability zones.
- D. Create multiple datastores for storage of VM.

Correct Answer: B

Community vote distribution

C (100%)

  **lurker8000** 4 months ago

Agree with C.

A replication factor of four would mean that each data block is replicated across four different nodes. While increasing the replication factor can improve fault tolerance, Cisco HyperFlex typically operates with a replication factor of three for high availability and performance.

upvoted 1 times

  **Rick099** 1 year, 4 months ago

Selected Answer: C

Logical Availability Zones are created to avoid multiple node and component failures on large clusters, and to increase cluster resiliency. Logical Availability Zones can only be enabled with clusters that are 8 converged nodes or larger.

https://intersight.com/help/appliance/features/hyperflex/configure#hyperflex_policies

upvoted 2 times

A Cisco UCS chassis with several blade servers, a pair of IOMs, and a pair of Cisco UCS Fabric Interconnects are installed in a data center.

Which two high-availability architectures are recommended between the IOMs and fabric interconnects? (Choose two.)

- A. One, two, four and eight links are supported to be used between each IOM and fabric interconnect.
- B. IOM A is connected to the fabric interconnect that is acting as fabric A, and IOM B is connected to the fabric interconnect that is acting as fabric B.
- C. Any combination of one to eight links is supported to be used between each IOM and fabric interconnect.
- D. The IOMs are connected with multiple straight and cross-connect links toward a pair of nonclustered fabric interconnects.
- E. Straight and cross-connect links should be used between the IOMs and fabric interconnects to offer higher availability.

Correct Answer: AB

Community vote distribution

AB (100%)

🗨️ 👤 **CoAsT_x** 2 months ago

Selected Answer: AB

It is A and B. lurker notes E, but cross connect links are not supported at all between IOMs and FIs. While A doesn't exactly mention availability, it does though. "used between each IOM and FI" infer that there is a connection between both IOMs and both FIs per chassis. UCS admin for 10 years. :)

upvoted 1 times

🗨️ 👤 **lurker8000** 4 months ago

I think it's B and E:

While the number of links (1, 2, 4, 8) is technically supported, this option does not emphasize high availability or resilient architecture between the IOMs and fabric interconnects. The number of links matters, but the high-availability architecture is more critical, which is described in options B and E.

upvoted 1 times

🗨️ 👤 **Rick099** 1 year, 4 months ago

Selected Answer: AB

The min # of links required to bring an I/O module or a chassis online are 1,2,4,8.

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/CLI-User-Guides/Infrastructure-Mgmt/3-2/b_UCSM_CLI_Infrastructure_Management_Guide_3_2/b_UCSM_CLI_Infrastructure_Management_Guide_3_2_chapter_011.html

upvoted 3 times

A company purchased a Cisco UCS system with four 5108 Chassis populated with UCS B-Series servers. The environment consists of nine Cisco HyperFlex rack servers. An engineer must provision virtual machines using an orchestration tool to optimize overall deployment time.

Which orchestration must be used to meet this requirement?

- A. Cisco Intersight
- B. Cisco DNA Center
- C. Cisco AppDynamic
- D. Cisco Network Assurance

Correct Answer: A

Currently there are no comments in this discussion, be the first to comment!

A company needs a solution to virtualize their legacy desktop computers. The solution must be highly resilient and able to withstand the loss of a data center site. Each compute node requires a minimum of 20 disk drives.

Which solution meets these requirements?

- A. thin client installed at each user's end station with a separate profile to retrieve sensitive data using the RDP protocol solution based on Cisco HX220 compute nodes
- B. data center site with a strict security policy that provides access to users using a bespoke software client solution based on Cisco HX240 compute nodes
- C. virtual desktop infrastructure solution installed in multiple data centers that are close to the end users solution based on Cisco HX240 compute nodes
- D. dedicated virtual machine per hypervisors for each user in a centralized data center solution based on Cisco HX220 compute nodes

Correct Answer: C

Currently there are no comments in this discussion, be the first to comment!

A network consultant must design a high availability interconnection of a clustered Cisco UCS Fabric Interconnect toward two upstream Layer 2 switches. The Ethernet interconnection must use all redundant links and have no impact on the STP domain size.

Which connectivity solution must be used?

- A. Cisco UCS Fabric Interconnect in switch mode with dual uplinks distributed evenly to the upstream switches
- B. Cisco UCS Fabric Interconnect in switch mode with a single uplink toward each upstream switch to eliminate any change to the STP domain size
- C. Cisco UCS Fabric Interconnect in end host mode with a single uplink toward each upstream switch to eliminate any change to the STP domain size
- D. Cisco UCS Fabric Interconnect in end host mode with dual uplinks distributed evenly to the upstream switches

Correct Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

An engineer is expanding an existing iSCSI environment and must choose an addressing scheme. When checking one of the devices, the engineer discovers this target device naming used:

```
{prefix}.2019-01.com.acme123:00.ff0cb21.bcb392817521
```

Which addressing scheme was used in the environment?

- A. Network Address Authority
- B. Extended Unique Identifier
- C. iSCSI Qualified Name
- D. Fibre Channel ID

Correct Answer: C

Community vote distribution

C (100%)

🗨️ **Rick099** 4 months ago

Selected Answer: C

Answer C. You can check the RFC 3720 / Section 3.2.6.3.1. Type "iqn." (iSCSI Qualified Name)
upvoted 2 times

🗨️ **mrkingsal** 10 months, 3 weeks ago

Selected Answer: C

I feel like this is an IQN
upvoted 2 times

A storage engineer requires a solution that achieves multiple paths between a server and storage array. The design must support complete traffic isolation and support a 50% growth in the next year.

Which FCoE solution meets these requirements?

- A. single switch with multiple links
- B. at least two VSANs across the fabric
- C. at least two FC zone sets with multiple zones
- D. single VSAN across the fabric


Correct Answer: *B*

 **Wasamela** 8 months ago

Correct.

upvoted 1 times

Refer to the exhibit.



A high-availability setup for iSCSI is designed with the host running iSCSI multipath software. Each IP storage port on the MDS switch exports the same two Fibre Channel target ports with different iSCSI target names.

In this design, how many iSCSI sessions are created from the host so that recovery occurs if any component fails?

- A. 2
- B. 3
- C. 4
- D. 5

Correct Answer: C

Currently there are no comments in this discussion, be the first to comment!

A network architect designs a high-availability architecture for the Cisco UCS Fabric Interconnects that are installed in a data center. The architecture must support clustering and redundancy.

Which architecture meets these requirements?

- A. dual connections from port L1 to L1 and one connection from port L1 to L1 with twinax cables
- B. dual connections from port L1 to L2 and one connection from port L2 to L1 with twinax cables
- C. dual connections from port L1 to L2 and one connection from port L2 to L1 with Ethernet cables
- D. dual connections from port L1 to L1 and one connection from port L2 to L2 with Ethernet cables

Correct Answer: D

Community vote distribution

C (67%)

D (33%)

🗨️ **lurker8000** 3 months, 2 weeks ago

Selected Answer: D

I'll go with D on this one based on the doc I attached.

upvoted 1 times

🗨️ **lurker8000** 3 months, 3 weeks ago

sorry I meant D. L1 to L1 and L2 to L2

upvoted 1 times

🗨️ **lurker8000** 3 months, 3 weeks ago

Did a bit more reading, I agree with C now based on this article:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/gui/config/guide/141/UCSM_GUI_Configuration_Guide_1_4_1_chapter4.html#:~:text=T

upvoted 1 times

🗨️ **lurker8000** 4 months ago

Why is twinax a good solution? Why not B?

upvoted 1 times

🗨️ **Wasamela** 8 months ago

Selected Answer: C

Answer should be C: dual connections between L1 and L2 ports on each fabric interconnect with standard Ethernet cables provides a redundant high-availability architecture in a data center.

upvoted 2 times

A company has added a new FCoE storage array to increase the performance of datastore access for end users. This storage array is directly connected to the Cisco Nexus 5672UP switches using four 10G port-channel interfaces. After deployment, the end users fail to access the datastore compared to the older solution. A network engineer verified the port-channel interface and finds that the congestion occurs when PFC frames are generated at the FCoE ingress interface. Which action must be taken to resolve the congestion issue?

- A. Reserve bandwidth for FCoE VLAN on the port-channel interface connected with storage.
- B. Enable the FCoE class in QoS on Cisco Nexus 5672UP switches.
- C. Increase the queuing buffer in the port-channel interface connected to the storage array.
- D. Add an additional FCoE VLAN to the SAN network.

Correct Answer: C

Community vote distribution

A (50%)

B (50%)

🗨️ **lurker8000** 3 months, 2 weeks ago

Selected Answer: B

Check comment below.

upvoted 1 times

🗨️ **lurker8000** 4 months ago

I will go for B:

The congestion occurs when PFC (Priority Flow Control) frames are generated at the FCoE ingress interface. In an FCoE (Fibre Channel over Ethernet) environment, PFC is used to prevent packet loss during congestion by pausing transmission on a specific priority class, but it requires proper Quality of Service (QoS) configuration to ensure that the correct traffic (in this case, FCoE traffic) is handled appropriately.

Enabling the FCoE class in QoS on the Cisco Nexus 5672UP switches ensures that FCoE traffic is prioritized and that PFC frames are processed properly. This avoids congestion by managing the flow of FCoE traffic more effectively, preventing bottlenecks when high traffic load occurs.

upvoted 1 times

🗨️ **Rick099** 1 year, 4 months ago

Selected Answer: A

If the egress FC port is congested, the switch sends PFC frames to the servers (storage array). The PFC frames are sent to reduce its FCoE rate and avoid a drop. So, Identify the source of the congestion and try to resolve it by increasing the FC bandwidth.

https://www.cisco.com/en/US/docs/switches/datacenter/nexus5500/sw/troubleshooting/guide/n5500_ts_fcoe.pdf

upvoted 1 times

A database administrator experiences delay when performing storage replication between primary and secondary data centers. The data centers are located 20 kilometers apart and are connected using a 500 Mbps link. The deployment was implemented using an FCIP tunnel and a pair of Cisco MDS 9250i Series Switches. The network engineer decided to enable QoS to prioritize replication traffic. Which QoS model must be used to resolve the performance issues?

- A. Resource Reservation Protocol
- B. weighted round robin
- C. differentiated services
- D. strict priority queuing

Correct Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **Wasamela** 8 months, 2 weeks ago

Selected Answer: C

Agree, answer is DiffServ

upvoted 1 times

🗨️ 👤 **Rick099** 1 year, 4 months ago

Selected Answer: C

- DSCP is the answer.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/6_2/configuration/guides/ip_services/nx-os/ipsvc/cfcip.html

upvoted 1 times

🗨️ 👤 **f143c37** 1 year, 8 months ago

Selected Answer: C

If the question in the exam is specifically asking for a "QoS Model" then it's DSCP and not strict priority queuing

upvoted 1 times

A customer must install a new Hyperflex Stretched cluster in two sites. Low latency and predictable high performance are the top priorities for this setup due to the extensive usage of data-intensive applications across the cluster. Which HyperFlex storage solution is recommended to meet these requirements?

- A. All-HDD
- B. Hybrid
- C. All-SAN
- D. All-Flash

Correct Answer: D

Community vote distribution

D (100%)



 **Wasamela** 8 months ago

Selected Answer: D

Correct

upvoted 1 times

The end users report issues with datastore reachability between the newly installed virtual machines (VMs) and the storage array. The VMs are deployed on a Cisco C-Series server, directly connected with Cisco Nexus 5672UP switches over FCoE VLAN. The data traffic on VLAN 99, which is designated as a native VLAN, reaches its default gateway, but FCoE VLAN 99 fails to access the datastore. Which action resolves the problem?

- A. Implement FCoE traffic on VLAN 10 and data traffic on VLAN 99.
- B. Implement host-facing FCoE ports as spanning-tree port type edge.
- C. Configure the FCoE VLAN in the VSAN database before including it in the trunk port.
- D. Configure the FCoE VLAN traffic on a separate interface from any other VLANs that traverse the network.

Correct Answer: A

Community vote distribution

C (67%)

A (33%)

🗨️ 👤 **OLENDE** 1 month, 1 week ago

Selected Answer: A

I think that FCoE is not supported on native vlan. So, the answer may be A

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/93x/san_switching/configuration/guide/b-cisco-nexus-9000-nx-os-san-switching-configuration-guide-933/b-cisco-nexus-9000-nx-os-san-switching-configuration-guide-933_chapter_01111.pdf

upvoted 1 times

🗨️ 👤 **lurker8000** 3 months, 3 weeks ago

Agree with C on this

upvoted 1 times

🗨️ 👤 **Wasamela** 8 months, 2 weeks ago

Selected Answer: C

I think that FCoE VLAN 99 needs to be configured in the VSAN database for the switch to recognize the traffic as FCoE

upvoted 2 times

A network administrator must create a redundant vHBA for the mission-critical virtual machines (VMs) installed on a Cisco UCS Fabric Interconnect blade server. The VMs access multiple datastores using VSAN55 and VSAN95. The datastore on which the VMs are hosted must be accessible if any of the uplinks goes offline. Which action meets these requirements?

- A. Attach VSAN55 to vHBA1 and VSAN95 to vHBA2.
- B. Assign VSAN55 and VSAN95 to vHBA1 and vHBA2.
- C. Enable failover for the vHBA with VSAN55 and VSAN95 enabled.
- D. Configure load-balancing for vHBAs with VSAN55 and VSAN95 attached.

Correct Answer: A

Community vote distribution

B (67%)

A (33%)

🗨️ 👤 **CoAsT_x** 2 months ago

Selected Answer: A

Cannot assign multiple VSANs to a vHBA. It is a single dropdown with one selection (unlike vNICs). Attaching one VSAN to one vHBA is the only option that is applicable here. Answer is A.

upvoted 1 times

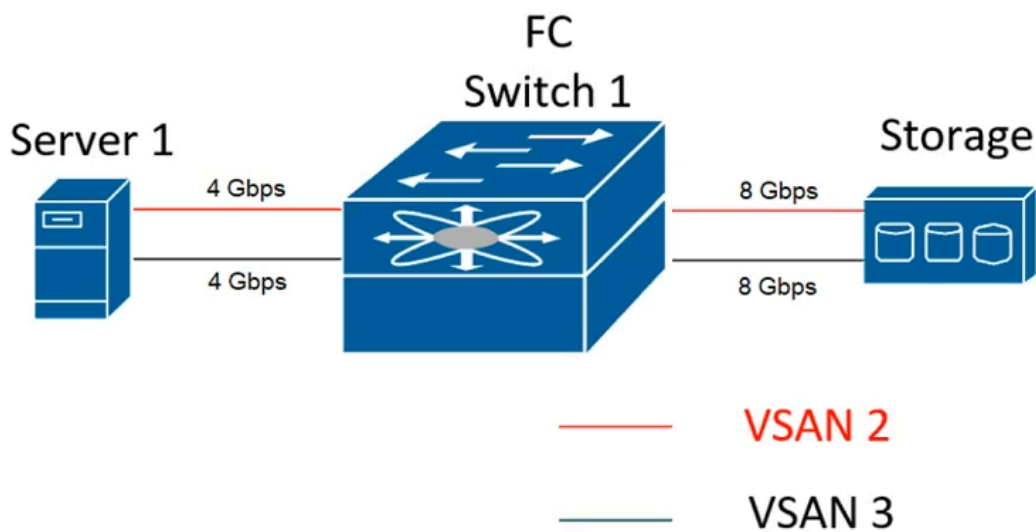
🗨️ 👤 **Wasamela** 8 months ago

Selected Answer: B

Answer should be B "assign" as you need to create multiple vHBAs and assign the different VSANs to them.

upvoted 2 times

Refer to the exhibit.



Server 1 fails to connect to the storage array over the Storage Area Network (SAN) of the Cisco MDS 9000 Series Switch. The requirements are to redesign the storage network and keep these considerations in mind:

- The traffic from each of the servers must be redundant and isolated.
- The design must tolerate hardware and software failures and upgrades of SAN fabric devices.

Which action must be taken to meet these requirements?

- Enable NPV and F-Port-channel on ports that face Server 1 on the Fibre Channel switch to create redundant paths.
- Create a SAN port channel that faces the storage device to sustain a link failure.
- Add an additional supervisor to the Fibre Channel switch to support nondisruptive upgrades.
- Place an additional Fibre Channel switch to create two physically independent storage fabrics.

Correct Answer: D

Community vote distribution

D (100%)

Wasamela 8 months, 2 weeks ago

Selected Answer: D

Agree, by creating two separate physical fabrics, failure in one of the fabrics won't disrupt communication on the second one.
upvoted 3 times

Rick099 1 year, 4 months ago

Selected Answer: D



Answer D.
upvoted 4 times

A network architect designs the iSCSI-based storage area network solution. The iSCSI datastore must be designed with high availability, minimal downtime, and no single point of failure. The requirement is to fail over I/O to alternate paths based on SCSI sense codes without reliance on network failures.

Which feature must be deployed to meet these requirements?

- A. hardware-based flow control
- B. port security
- C. iSCSI multipath
- D. NIC teaming

Correct Answer: C

  **lurker8000** 3 months, 3 weeks ago

Correct, C

upvoted 1 times

A network consultant reviews the requirements for the Storage Area Network (SAN) design. The proposed solution must offer a very efficient port usage and eliminate the need for a separate physical layer of switches and their associated ISLs. The design is not expected to require future expansion.

Which SAN topology must be used to meet these requirements?

- A. core-edge
- B. collapsed-core
- C. core-distribution-access
- D. edge-core-edge

Correct Answer: B

Community vote distribution

B (100%)



 **Wasamela** 8 months ago

Selected Answer: B

Correct

upvoted 2 times