



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

What is an accurate statement about Governance?

- A. Restrict access to certain users
- B. Ensure removal of security risks
- C. Authorize policy-making decisions
- D. Demonstrate adherence to policies

Suggested Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **Saeid** Highly Voted 3 years, 9 months ago

Passed my exam, about 60% of my questions came from this dump. Had so many new questions. I will try to write it here as much as I remember:

- 1- which component is responsible for cache management? - controller
 - 2- Problem management: Prevents incidents that share common symptoms or root causes from reoccurring and minimizes the adverse impact of incidents that cannot be prevented.
 - 3- advantage of cumulative backup over incremental? shorter restore time
 - 4- VPN Definition
 - 5- Dell EMC VxRack SDDC: delivers the simplest path to hybrid cloud with an automated elastic cloud infrastructure at rack scale.
 - 6- "front-end controller" has processing logic that executes the appropriate transport protocol, such as Fibre Channel, iSCSI, FICON, or FCoE for storage connections.
 - 7- Common Internet File System / Server Message Block (CIFS/SMB)
- upvoted 8 times

🗨️ 👤 **fer162** Highly Voted 4 years, 4 months ago

Correct, from participant guide "Governance is the authority for making policies such as defining access rights to users based on their roles and privileges"

upvoted 5 times

🗨️ 👤 **SupremeND** Most Recent 2 years, 6 months ago

Selected Answer: C

Governance is the authority for making policies such as defining access rights to users based on their roles and privileges. Risk management involves identifying resources that should not be accessed by certain users in order to preserve confidentiality, integrity, and availability. In this example, compliance management assures that the policies are being enforced by implementing controls such as firewalls and identify management systems.

upvoted 1 times

🗨️ 👤 **whickz22** 3 years, 1 month ago

Selected Answer: C

Pass my Exam, Mostly this is still valid, and as saeid mention earlier, there is a plenty new & modified question, but still can made it, goodluck everyone !

upvoted 1 times

🗨️ 👤 **jmcastro** 2 years, 8 months ago

Hi whickz22, which version of ISM did you pass?

upvoted 1 times

🗨️ 👤 **skyrn** 3 years, 5 months ago

Could you please send me a copy of the ISMv4 pdf Guide.Thanks.

upvoted 1 times

🗨️ 👤 **jeremy15n** 3 years, 7 months ago

If you need the ISMv4 Participant Guide add me on Discord and ill send it to you. Jay-#5335

upvoted 2 times

🗨️ 👤 **whickz22** 3 years, 2 months ago

hi jay, already sent a friend req on discord, kindly accept, thanks
upvoted 1 times

🗨️ 👤 **Saeid** 3 years, 9 months ago

other questions I remember from the exam:

8- policy engine automated moving data from primary storage to archiving storage.

9- What is the main purpose of creating local replica in a data center?

options: act as an archiving source or restart the business

10- Compliance

11- vRealize operation

12- if customer has access to storage, operating systems, and network, what does it called? IAAS or PAAS?

upvoted 4 times

🗨️ 👤 **shaolong1231** 3 years, 10 months ago

Pass the exam, about 80% of the exam questions come from this dump (in Taiwan), and the answers need to be corrected according to "Discussion"

upvoted 1 times

🗨️ 👤 **Pink_Leche** 4 years, 1 month ago

learner, have you found anything more accurate ?

upvoted 1 times

🗨️ 👤 **Poopydoo** 4 years, 4 months ago

@learner2, have you pass the exam already, is there a new one now?

upvoted 1 times

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

i failed, looking for new.. i see some 180QAS

Don't know if its valid..can't find the 180QAs to see and check.

upvoted 2 times

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

What was your score?

upvoted 1 times

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

How many items here are in new exam?

upvoted 1 times

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

<https://www.passleader.com/dea-1tt4.html>

upvoted 2 times

🗨️ 👤 **Batuhan** 4 years, 4 months ago

Is this valid or no more?I'm confused by learner2

upvoted 1 times

🗨️ 👤 **learner** 4 years, 4 months ago

does any one have the updated ISMv4, this no more valid

upvoted 2 times

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


please send me the ISMv4 guide you have

upvoted 1 times

Which exhibit represents a Type 1 hypervisor?


A.

Exhibit 1:




B.

Exhibit 2:




C.

Exhibit 3:



D.

Exhibit 4:



Suggested Answer: A

SupremeND 2 years, 6 months ago

A is a Type 1 hypervisor

B is a Type 2 hypervisor

upvoted 1 times

gamalaa 2 years, 4 months ago

C is Type 2 not B

upvoted 1 times

[Removed] 3 years, 3 months ago



A is correct

upvoted 2 times

Pink_Leche 3 years, 7 months ago

The question on my exam asked for the Type 2 Hypervisor.

upvoted 3 times

  **inf_ph** 3 years, 3 months ago

Hey Pink_Leche,

When did you take the exam? How was it? Did this dump help you?

upvoted 1 times

An Ethernet switch has eight ports and all ports must be operational for 24 hours, from Monday through Friday. However, a failure of Port 4 occurs as follows:

"ç Monday = 8 PM to 9 PM

"ç Wednesday 7 PM to 9 PM

"ç Thursday = 6 AM to 9 AM

"ç Saturday = 4 PM to 7 PM

What is the MTTR of Port 4?

A. 2 hours

B. 3 hours

C. 28 hours

D. 38 hours

Suggested Answer: A

Community vote distribution

A (100%)

🗨️ 👤 **Varunraj** Highly Voted 4 years, 3 months ago

No need to consider the downtime for Saturday since the question mentioned about Monday to Friday. So MTTR will be 6h downtime/3 incidents = 2

upvoted 18 times

🗨️ 👤 **SupremeND** Most Recent 2 years, 6 months ago

Selected Answer: A

MTTR=total time to repair/number of incidents=6/3=2

upvoted 1 times

🗨️ 👤 **zahidkhader** 3 years, 1 month ago

Selected Answer: A

Monday to Friday. So MTTR will be 6h downtime/3 incidents = 2

upvoted 1 times

🗨️ 👤 **Mahrizou** 3 years, 10 months ago

MTBF: Average time available for a system or component to perform its normal operations between failures

▮ MTBF= Total uptime / Number of failures

MTTR: Average time required to repair a failed component

▮ MTTR = Total downtime / Number of failures

upvoted 2 times

🗨️ 👤 **madmax67** 4 years, 4 months ago

MTTR = Meantime to Repair.

ISMv4 PArticipant Guide: MTTR: Average time required to repair a failed component

MTTR = Total downtime / Number of failures

Given downtimes per failure:

8 to 9 PM = 1h

7 to 9 PM = 2h

6 to 9 AM = 3h

4 to 7 PM = 3h

Calculation: 9h downtime / 4 incidents = 2,25h MTTR

Conclusion: 2 h 15 m should be the correct answer.

upvoted 3 times

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

Downtime between Mon - Fri only. So only 6 hours /3 failures = 2 hours

upvoted 4 times

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

Kindly send me a copy of the ISMv4 Participant Guide.Thanks.

upvoted 1 times

🗨️ 👤 **hizohizo15** 3 years ago

its monday to friday...so only 6hrs downtime..dont include saturday..

upvoted 1 times

What accurately describes an iSCSI HBA?

- A. Performs both iSCSI and TCP/IP processing
- B. Requires CPU cycles for TCP/IP processing
- C. Requires a software adapter with built-in functionality
- D. Enables transmission of data through iSCSI tunneling

Suggested Answer: A

Reference:

<https://www.dell.com/downloads/global/power/ps4q09-20090419-McConnell.pdf>

Community vote distribution

A (100%)

🗨️ 👤 **SupremeND** 2 years, 6 months ago

Selected Answer: A

- iSCSI HBA
 - Performs both iSCSI and TCP/IP processing
 - Frees-up CPU cycles of compute system for business applications
- upvoted 1 times

🗨️ 👤 **suzcz** 3 years, 9 months ago

A is correct. iSCSI initiator HBA performs both iSCSI & TCP IP

upvoted 1 times

🗨️ 👤 **Mahrizou** 3 years, 10 months ago

- ☒ iSCSI HBA
 - ☒ Performs both iSCSI and TCP/IP processing
 - ☒ Frees-up CPU cycles of compute system for business applications
- upvoted 4 times



When using the erasure coding technique in an object-based storage system, data is divided into 12 data segments and 4 coding segments. What is the maximum number of disk drive failures against which data is protected by the configuration?

- A. 3
- B. 4
- C. 8
- D. 12

Suggested Answer: A

Community vote distribution

B (100%)

  **madmax67** Highly Voted  4 years, 5 months ago

In a typical erasure coded storage system, a set of n disks is divided into m disks to hold data and k disks to hold coding information, where n , m , and k are integers. The coding information is calculated from the data. If up to k of the n disks fail, their contents can be recomputed from the surviving disks.

In this example $k=4$, so the correct answer is B.

upvoted 17 times

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

B. 4

"Provides space-optimal data redundancy to protect data loss against multiple drive failures

- A set of n disks is divided into m disks to hold data and k disks to hold coding information
- Coding information is calculated from data

The figure illustrates an example of dividing a data into nine data segments ($m = 9$) and three coding fragments ($k = 3$). The maximum number of drive failure supported in this example is three."

upvoted 8 times



  **[Removed]** Most Recent  5 months, 3 weeks ago

Selected Answer: B

In erasure coding, data is divided into data segments and coding segments. The total number of segments is the sum of data segments and coding segments. In this configuration, there are 12 data segments and 4 coding segments, making a total of 16 segments.

The maximum number of disk drive failures the system can tolerate is equal to the number of coding segments, which is 4 in this case. This means that data can still be reconstructed even if up to 4 disks fail simultaneously.

upvoted 1 times



  **SupremeND** 2 years, 6 months ago

Selected Answer: B

The figure illustrates an example of dividing a data into nine data segments ($m = 9$) and three coding fragments ($k = 3$). The maximum number of drive failure supported in this example is three.

Erasure coding offers higher fault tolerance (tolerates k faults) than replication with less storage cost.

upvoted 1 times

  **KimZhou** 2 years, 10 months ago



Selected Answer: B

"Provides space-optimal data redundancy to protect data loss against multiple drive failures

- A set of n disks is divided into m disks to hold data and k disks to hold coding information
- Coding information is calculated from data

The figure illustrates an example of dividing a data into nine data segments ($m = 9$) and three coding fragments ($k = 3$). The maximum number of drive failure supported in this example is three."

upvoted 1 times

  **amouraa** 2 years, 11 months ago

Selected Answer: B

B 4 , not A

upvoted 1 times

🗨️ **Albert98llioni** 3 years, 3 months ago

The answer is 3. $12/4=3$

upvoted 1 times

🗨️ **Salah94** 3 years, 3 months ago

Selected Answer: B

Can tolerate up to K disk (here $k=4$)

upvoted 2 times

🗨️ **inf_ph** 3 years, 3 months ago

Hey @Salah94

Have you taken the exam?

upvoted 1 times

🗨️ **Ivan_Miran** 3 years, 4 months ago

Correct answer is B

<https://searchstorage.techtargget.com/definition/erasure-coding>

upvoted 1 times

🗨️ **Jhonas39** 3 years, 6 months ago

correct A

$12/4= 3$

upvoted 2 times

🗨️ **Jhonas39** 3 years, 6 months ago

I think the correct is A.

Because the data segments is 12 and the coding segments is 4 then:

$M/K=9/4= 3$

upvoted 1 times

What is a key advantage of cloning a virtual machine (VM)?

- A. Ability to run new VMs simultaneously with the same IP address.
- B. Reduces the time required to create multiple copies of a VM.
- C. Improves the performance of an existing VM and new VMs.
- D. Eliminates the need to create multiple snapshots for a VM.

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ **SupremeND** 2 years, 6 months ago

Selected Answer: B

In general, installing a guest operating system and applications on a VM is a time consuming task. With clones, administrators can make many copies of a virtual machine from a single installation and configuration process.

upvoted 1 times

🗨️ **nyashac** 2 years, 7 months ago

Selected Answer: B

With cloning, you can quickly deploy multiple virtual machines with the same setup.

<https://www.liquidweb.com/kb/vmware-cloning/#:~:text=With%20cloning%2C%20you%20can%20quickly,in%20just%20a%20few%20clicks.>

upvoted 1 times

🗨️ **andikovaci** 3 years, 1 month ago

Selected Answer: B

correct

upvoted 1 times

🗨️ **andikovaci** 3 years, 1 month ago

Correct

upvoted 1 times

🗨️ **gilbertoinf** 3 years, 2 months ago

I'm confused with this question. Because that option B tell about "reduce time to create multiples copies of vms". The clone just do a copy of VM, but it don't reduce the time to make this. For me, the option C make more sense.

upvoted 1 times

🗨️ **gilbertoinf** 3 years, 2 months ago

Not option C, but D

upvoted 1 times

🗨️ **jefflee78** 3 years, 3 months ago

Correct

upvoted 1 times

A company is using an intrusion detection and prevention system (IDPS) to scan and analyze intrusion events. Which IDPS technique uses a database that contains known attack patterns?

- A. Role-based
- B. Profile-based
- C. Signature-based
- D. Anomaly-based

Suggested Answer: C

Reference:

<https://pdfs.semanticscholar.org/e874/32283e90da5bceb3b17550570ec3a210f43c.pdf>

Community vote distribution

C (100%)

🗨️ **SupremeND** 2 years, 6 months ago

Selected Answer: C

- Signature-based detection technique:
 - Scans for signatures to detect an intrusion
 - Effective only for known threats
- upvoted 1 times

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

The answer is anomaly-based as they have mentioned patterns in the question

upvoted 1 times

🗨️ **salvo92** 3 years, 5 months ago

Correct

upvoted 2 times

🗨️ **Dazzie** 3 years, 6 months ago

Correct.

upvoted 2 times

DRAG DROP -

In a backup operation, what are the correct sequence of steps after the backup server initiates the scheduled backup process?

Select and Place:

Storage node sends data to the backup device	STEP 1
Backup server instructs the storage node to load the backup media in the backup device	STEP 2
Backup server updates the backup catalog	STEP 3
Backup clients send the data to the storage node and updates the backup catalog on the backup server	STEP 4
Storage node sends metadata and media information to the backup server	STEP 5

Suggested Answer:

Storage node sends data to the backup device	Backup server instructs the storage node to load the backup media in the backup device
Backup server instructs the storage node to load the backup media in the backup device	Backup clients send the data to the storage node and updates the backup catalog on the backup server
Backup server updates the backup catalog	Storage node sends data to the backup device
Backup clients send the data to the storage node and updates the backup catalog on the backup server	Storage node sends metadata and media information to the backup server
Storage node sends metadata and media information to the backup server	Backup server updates the backup catalog

 **Maryjn3** 3 years, 1 month ago

- Backup server instructs the storage node to load the backup media in the backup device
 - Backup client send data to the storage node and updates the backup catalog on the backup server
 - Storage node sends data to the backup device
 - Storage node sends metadata and media information to the backup server
 - Backup server updates the backup catalog
- upvoted 1 times

 **Ivan_Miran** 3 years, 3 months ago

"The backup server retrieves the backup-related information from the backup catalog and, based on this information, (1) instructs the storage node to load the appropriate backup media into the backup devices. (2)/(3) Simultaneously, it instructs the backup clients to gather the data to be backed up and send it over the network to the assigned storage node. After the backup data is sent to the storage node, the client sends some backup metadata (the number of files, name of the files, storage node details, and so on) to the backup server. The storage node receives the client data, organizes it, and sends it to the backup device. (4) The storage node then sends additional backup metadata (location of the data on the backup device, time of backup, and so on) to the backup server. (5) The backup server updates the backup catalog with this information.

Information Storage and Management. Wiley. Kindle Edition.

upvoted 1 times

 **mozgachev** 3 years, 5 months ago

- (1) Backup server initiates scheduled backup process.
- (2) Backup server retrieves backup-related information from the backup catalog.

(3a) Backup server instructs storage node to load backup media in the backup device.

(3b) Backup server instructs backup clients to send data to be backed up to the storage node.

(4) Backup clients send data to storage node and update the backup catalog on the backup server.

(5) Storage node sends data to the backup device

(6) Storage node sends metadata and media information to the backup server

(7) Backup server updates the backup catalog

upvoted 2 times

  **mozgachev** 3 years, 5 months ago

correct

upvoted 1 times

In an archiving environment, how is a content address created in a content addressable storage (CAS) system to store an object?

- A. From the metadata of an object.
- B. By using the logical address of the memory space.
- C. From the binary representation of an object.
- D. By using the physical address of the storage device.

Suggested Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **SupremeND** 2 years, 6 months ago

Selected Answer: C

CAS stores user data and its attributes as an object. The stored object is assigned a globally unique address, which is known as a content address (CA). This address is derived from the binary representation of an object. Content addressing eliminates the need for application servers to understand and manage the physical location of objects on a storage system.

upvoted 1 times

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

C

"CAS is an object-based storage device designed for secure online storage and retrieval of fixed content. CAS stores user data and its attributes as an object. The stored object is assigned a globally unique address, known as a content address (CA). <<This address is derived from the object's binary representation.>>

Information Storage and Management . Wiley. Kindle Edition.

upvoted 2 times

🗨️ 👤 **mathmad** 4 years, 6 months ago

should be C

upvoted 4 times

🗨️ 👤 **amtavt** 4 years, 5 months ago

Yes, it's C

"CAS stores user data and its attributes as an object. The stored object is assigned a globally unique address, which is known as a content address (CA). This address is derived from the binary representation of an object. Content addressing eliminates the need for application servers to understand and manage the physical location of objects on a storage system."

upvoted 4 times

What is a benefit of link aggregation?

- A. Optimizes fabric performance by distributing network traffic across ISLs.
- B. Improves ISL utilization by increasing the network traffic on a specific ISL.
- C. Yields a higher throughput by combining two parallel ISLs into a single physical ISL.
- D. Saves bandwidth by restricting traffic to a specific potentially congested ISL.

Suggested Answer: A

Reference:

<https://www.mycloudwiki.com/san/fc-san-topologies/>

🗨️ **salmoallemi** Highly Voted 3 years, 9 months ago

A is correct, C is not correct that it should be combined into a logical ISL not physical.
upvoted 7 times

🗨️ **EslamYaseen** Most Recent 3 years, 7 months ago

A is the correct answer
as the question asks for the benefits of Link aggregation
C Describes the definition of link aggregation
upvoted 1 times

🗨️ **jan2020** 3 years, 11 months ago

Link aggregation combines two or more parallel ISLs into a single logical ISL.
"Two or more" not only two,so A correct.
upvoted 1 times

🗨️ **ahmet196621** 3 years, 11 months ago

C is not correct as "combining two parallel ISLs" does NOT put it "into a single PHYSICAL ISL". The word "physical" is completely wrong here. So i think the correct answer is A.
upvoted 3 times

🗨️ **Jiki_Jiki** 4 years ago

Both A & C are correct.
Link aggregation combines two or more parallel ISLs into a single logical ISL, called a port-channel, yielding higher throughput than a single ISL could provide.
For example, the aggregation of 10 ISLs into a single port-channel provides up to 160 Gb/s throughput assuming the bandwidth of an ISL is 16 Gb/s. Link aggregation optimizes fabric performance by distributing network traffic across the shared bandwidth of all the ISLs in a port-channel. This allows the network traffic for a pair of node ports to flow through all the available ISLs in the port-channel rather than restricting the traffic to a specific, potentially congested ISL.
upvoted 1 times

🗨️ **Kate321** 3 years, 11 months ago

C is not correct beacuse is says physical and should be logical
upvoted 3 times

🗨️ **5kyFx** 4 years ago

The link is good, reading the content leads to answer C
upvoted 1 times

🗨️ **WhitWolf83** 4 years ago

Should be C based on the below.

"Link aggregation combines two or more parallel ISLs into a single logical ISL, called a port-channel, yielding higher throughput than a single ISL could provide"
upvoted 3 times

🗨️ **shav007** 3 years, 4 months ago

READ THE ANSWER AGAIN...IT SAYS PHYSICAL NOT LOGICAL

upvoted 1 times

  **WhitWolf83** 4 years ago

Should be C

upvoted 3 times

In a compute-based storage area network (SAN), what is a function of the metadata manager?

- A. Holds cluster-wide mapping information and load balancing in a compute-based SAN.
- B. Performs the I/O operations on the local storage of a compute system as requested by a client.
- C. Provides an interface to the software-defined storage controller.
- D. Controls all data read and write operations in a compute-based SAN.

Suggested Answer: A

 **fer162** Highly Voted 4 years, 4 months ago

"A compute-based SAN's control component, which is known as the metadata manager, serves as the monitoring and configuration agent. It holds cluster-wide mapping information and monitors capacity, performance, and load balancing."

upvoted 12 times

Why is it important for organizations to store, protect, and manage their data?

- A. Reduce the amount of data to be replicated, migrated, and backed up.
- B. Rapidly develop and deploy modern applications for business improvement.
- C. Develop new business opportunities by leveraging existing data.
- D. Reduce the complexity in managing the data center environment.

Suggested Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **SupremeND** 2 years, 6 months ago

Selected Answer: C

Data is the lifeblood of a rapidly growing digital existence, opening up new opportunities for businesses and gain a competitive edge. For example, an online retailer may need to identify the preferred product types and brands of customers by analyzing their search, browsing, and purchase patterns. This information helps the retailer to maintain a sufficient inventory of popular products, and also advertise relevant products to the existing and potential customers. It is essential for organizations to store, protect, process, and manage information in an efficient and cost-effective manner. Legal, regulatory, and contractual obligations regarding the availability, retention, and protection of data further add to the challenges of storing and managing information.

upvoted 1 times

🗨️ 👤 **mozgachev** 3 years, 5 months ago

correct!

"Data is the lifeblood of a rapidly growing digital existence, opening up new opportunities for businesses and gain a competitive edge. For example, an online retailer may need to identify the preferred product types and brands of customers by analyzing their search, browsing, and purchase patterns. This information helps the retailer to maintain a sufficient inventory of popular products, and also advertise relevant products to the existing and potential customers. It is essential for organizations to store, protect, process, and manage information in an efficient and cost-effective manner. Legal, regulatory, and contractual obligations regarding the availability, retention, and protection of data further add to the challenges of storing and managing information."

upvoted 2 times

A Fibre Channel switch has eight ports and all ports must be operational 24×7. However, a failure of Port 5 occurs on a particular week as follows:

☞ Monday = 11 AM to 12 PM

☞ Thursday = 7 AM to 11 AM

☞ Saturday = 11 AM to 1 PM

What is the availability of Port 5 in that week?

- A. 93%
- B. 94.1%
- C. 95.8%
- D. 97%

Suggested Answer: C

🗨️ **ssajith** Highly Voted 4 years, 1 month ago

Total hours in a day = 24

Total hours in a week = $24 \times 7 = 168$

Total hours ports are down in the week = $1 + 4 + 2 = 7$ hrs

Total expected availability = 100 %

Total % of unavailability = $7 \times 100 / 168 = 4.166667$

Availability of port = $100 - 4.2 = 95.8$ %

upvoted 15 times

🗨️ **KZM** Most Recent 3 years, 7 months ago

$\{(168-7)/168\} \times 100 = 95.8\%$

upvoted 4 times

🗨️ **madmax67** 4 years, 5 months ago

What is meant by 12 PM? Is it Noon or is it Midnight?

upvoted 3 times

🗨️ **Tigrediablo100** 4 years, 5 months ago

It's noon, the Midnight would 00:00 or 12 AM

upvoted 2 times

🗨️ **atidd** 4 years, 5 months ago

11 AM to 12 PM = 11 AM to Noon (midday) = 1 hr.

midnight is 12 AM.

upvoted 2 times

🗨️ **Dazzie** 3 years, 6 months ago

Noon bruh.

upvoted 1 times

Which layer of the OSI model provides the functional and procedural methods to transfer data between network entities?

- A. Data link
- B. Presentation
- C. Session
- D. Application

Suggested Answer: A

Reference:

https://en.wikibooks.org/wiki/Network_Plus_Certification/Management/OSI_Model

🗨️ 👤 **AdhamELHariry** 3 years, 3 months ago

Data Link

The data link layer provides the functional and procedural means to transfer data between network entities and to detect and possibly correct errors that may occur in the physical layer

upvoted 1 times

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

correct

upvoted 1 times

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

its D, Application Layer.

upvoted 1 times

🗨️ 👤 **rachid1761960** 3 years, 6 months ago

1. Physical Layer - Defines the electrical and physical specifications for devices.
2. Data Link Layer - Provides the functional and procedural means to transfer data between network entities. It also detects and possibly correct errors that may occur in the Physical Layer.
3. Network Layer - Transfers variable length data sequences from a source to destination through one or more networks while also maintaining a quality of service requested by the Transport Layer.
4. Transport Layer - Provides transparent transfer of data between end users, providing reliable data transfer services to the upper layers.
5. Session Layer - Controls the connections between computers. It establishes, manages, and terminates the connections between the local and remote application.
6. Presentation Layer - Establishes a context between the Application layer entities in which the high-layer entities can use different syntax and semantics.
7. Application Layer - Provides a user interface that enables user to access the network and applications.

upvoted 3 times

🗨️ 👤 **kens** 4 years, 1 month ago

https://en.wikibooks.org/wiki/Network_Plus_Certification/Management/OSI_Model

upvoted 1 times

What is an accurate statement about variable prefetching?

- A. Many blocks that have recently been written to storage by the compute system are retrieved and placed in cache.
- B. Many non-contiguous blocks that are not currently required can be retrieved from cache in advance.
- C. Storage system prefetches data in multiples of the original I/O request size from the compute system.
- D. Storage system prefetches many non-contiguous blocks of data, with varied sizes, from cache.

Suggested Answer: D

Community vote distribution

C (100%)

JBONINO Highly Voted 4 years ago

- A. Many blocks that have recently been written to storage by the compute system are retrieved and placed in cache. - NO// WRITE, MUST BE READ
 - B. Many non-contiguous blocks that are not currently required can be retrieved from cache in advance. - NO // Must to be contiguous
 - C. Storage system prefetches data in multiples of the original I/O request size from the compute system. - MOST ACCURATE - SO C for ME
 - D. Storage system prefetches many non-contiguous blocks of data, with varied sizes, from cach - NO // Must to be contiguous
- upvoted 8 times

Badr0093 Most Recent 5 months, 3 weeks ago

Correct answer is C
upvoted 1 times

NaelOsama 2 years, 10 months ago

Selected Answer: C

Prefetch: A prefetch or read-ahead algorithm is used when read requests are sequential. In a sequential read request, a contiguous set of associated blocks is retrieved. Several other blocks that have not yet been requested by the compute system can be read from the storage and placed into cache in advance. When the compute system subsequently requests these blocks, the read operations will be read hits. This process significantly improves the response time experienced by the compute system

upvoted 2 times

Ivan_Miran 3 years, 3 months ago

Must be "C"

"The intelligent storage system offers fixed and variable prefetch sizes. In fixed prefetch, the intelligent storage system prefetches a fixed amount of data. It is most suitable when host I/O sizes are uniform. << In variable prefetch, the storage system prefetches an amount of data in "multiples of the size of the host request". >> Maximum prefetch limits the number of data blocks that can be prefetched to prevent the disks from being rendered busy with prefetch at the expense of other I/Os.

Information Storage and Management . Wiley. Kindle Edition.

upvoted 1 times

mozgachev 3 years, 5 months ago

From ISM v4 -> Prefetch: A prefetch or read-ahead algorithm is used when read requests are sequential. In a sequential read request, a contiguous set of associated blocks is retrieved. Several other blocks that have not yet been requested by the compute system can be read from the storage and placed into cache in advance. When the compute system subsequently requests these blocks, the read operations will be read hits. This process significantly improves the response time experienced by the compute system.

upvoted 1 times

walther32 3 years, 5 months ago

<https://forum.huawei.com/enterprise/en/corpus-12891.html>

1. After receiving a data read request, the storage system prefetches data to the cache based on a multiple of the read length specified in the I/O request.

2. This policy applies to sequential read applications with variable-size data or to multi-user concurrent read applications with an unknown data prefetch amount, for example, multi-user Video on Demand (VOD) initiated by users at different bit rates.

C is correct.

upvoted 3 times

🗨️ 👤 **Tyrion0097** 3 years, 6 months ago

A prefetch or read-ahead algorithm is used when read requests are sequential. In a sequential read request, a contiguous set of associated blocks is retrieved. Several other blocks that have not yet been requested by the compute system can be read from the storage and placed into cache in advance. When the compute system subsequently requests these blocks, the read operations will be read hits. This process significantly improves the response time experienced by the compute system.

- Since it's a variable prefetch then D is correct.

upvoted 1 times

🗨️ 👤 **sheym** 4 years, 1 month ago

<https://www.dell.com/downloads/global/power/ps2q04-022.pdf> page 97

upvoted 2 times

🗨️ 👤 **kens** 4 years, 1 month ago

Cache prefetching

The EMC read-prediction algorithm—also known as read-ahead caching or prefetching—is adaptive in nature. The SP will prefetch data (assuming prefetching is enabled for the corresponding LUN) and fill the read cache only after two sequential reads that share spatial locality. By using this algorithm, the SP assumes that if there were two reads from the same location, there is a high probability that it will need that data from the next sequential location again. The SP will first check the read cache for the data. If the data is not present, it will then check the disk. This algorithm can help reduce SP-to-LUN latency by filling the read cache with prefetched data before it is actually needed.

upvoted 2 times

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

D. In variable prefetch, the storage system prefetches an amount of data in multiples of the size of the host request.

upvoted 2 times

🗨️ 👤 **Jamestyler22** 4 years, 4 months ago

It's B - Several other blocks that have not yet been requested by the compute system can be read from the storage and placed into cache in advance

upvoted 3 times

🗨️ 👤 **farquar** 4 years, 1 month ago

but B says read from cache not written to cache so it's wrong, C is correct.

upvoted 2 times

🗨️ 👤 **hakot** 4 years, 5 months ago

I actually think it is B, or maybe D. But maybe I am understanding this question wrong but here is what I found:

Prefetch: A prefetch or read-ahead algorithm is used when read requests are sequential. In a sequential read request, a contiguous set of associated blocks is retrieved. Several other blocks that have not yet been requested by the compute system can be read from the storage and placed into cache in advance. When the compute system subsequently requests these blocks, the read operations will be read hits. This process significantly improves the response time experienced by the compute system.

upvoted 1 times

🗨️ 👤 **mathmad** 4 years, 7 months ago

should be C

upvoted 3 times

DRAG DROP -

Match the Fibre Channel (FC) layers with their respective functions.

Select and Place:

FC Layers

FC-1
FC-2
FC-3
FC-4

FC Functions

Routing
Mapping interface
Encode/Decode
Common services

Suggested Answer:

FC Layers

FC-1
FC-2
FC-3
FC-4

FC Functions

FC-2
FC-4
FC-1
FC-3

Reference:

[https://www.google.com/url?](https://www.google.com/url?sa=t&rc=t=j&q=&esrc=s&source=web&cd=10&ved=2ahUKewjoirWjjKblAhXyxoUKHbl_AU4QFjAJegQIAxAC&url=https%3A%2F%2Fwww.dell.com%2Fcommunity%2Fs%2Fvjauj58549%2Fattachments%2Fvjauj58549%2Fconnectrix%2F4498%2F1%2FSC_Distance_Reference_GA-SG-061-00%255B1%255D.pdf&usg=AOvVaw3SqLy5ETVbbmIBfl579cca)

[sa=t&rc=t=j&q=&esrc=s&source=web&cd=10&ved=2ahUKewjoirWjjKblAhXyxoUKHbl_AU4QFjAJegQIAxAC&url=https%3A%2F%2Fwww.dell.com%2Fcommunity%2Fs%2Fvjauj58549%2Fattachments%2Fvjauj58549%2Fconnectrix%2F4498%2F1%2FSC_Distance_Reference_GA-SG-061-00%255B1%255D.pdf&usg=AOvVaw3SqLy5ETVbbmIBfl579cca](https://www.dell.com%2Fcommunity%2Fs%2Fvjauj58549%2Fattachments%2Fvjauj58549%2Fconnectrix%2F4498%2F1%2FSC_Distance_Reference_GA-SG-061-00%255B1%255D.pdf&usg=AOvVaw3SqLy5ETVbbmIBfl579cca)

 **hackidies** Highly Voted 4 years ago

Should :

FC0-physical


FC1-encode/decode

FC2-routing, flowcontrol

FC3-Common service

FC4-mapping

<https://www.packetflow.co.uk/fibrechannel-sanprotocolsexplained/>
upvoted 11 times

 **Pink_Leche** Most Recent 3 years, 7 months ago

My exam included FC-0

upvoted 2 times

 **Saeid** 3 years, 9 months ago

Page 306

FC-4 Mapping interface

FC-3 Common services

FC-2 Routing

FC-1 Encode/decode

upvoted 2 times

Which Dell EMC product is a hyper-scale storage infrastructure that provides universal accessibility with support for object and HDFS?

- A. ViPR Controller
- B. PowerMax
- C. Isilon
- D. Elastic Cloud Storage

Suggested Answer: D

Reference:

[https://www.google.com/url?](https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=10&ved=2ahUKewibs5yDjablAhVCzBoKHbc6AhgQFjAJegQIAxAC&url=https%3A%2F%2Fwww.dell.com%2Fcommunity%2Fs%2Fvjauj58549%2Fattachments%2Fvjauj58549%2Fsolutions-ch%2F516%2F1%2Fh13999-hadoop-ecs-data-services-wp.pdf&usg=AOvVaw1jTgO2qHQLxapR5zvgXWpH)

[sa=t&rct=j&q=&esrc=s&source=web&cd=10&ved=2ahUKewibs5yDjablAhVCzBoKHbc6AhgQFjAJegQIAxAC&url=https%3A%2F%2Fwww.dell.com%2Fcommunity%2Fs%2Fvjauj58549%2Fattachments%2Fvjauj58549%2Fsolutions-ch%2F516%2F1%2Fh13999-hadoop-ecs-data-services-wp.pdf&usg=AOvVaw1jTgO2qHQLxapR5zvgXWpH](https://www.dell.com/community/s/vjauj58549/attachments/vjauj58549/solutions-ch/516/1/h13999-hadoop-ecs-data-services-wp.pdf)


 **mozgachev** 3 years, 5 months ago

Dell EMC ECS

Provides a hyper-scale storage infrastructure that is specifically designed to support modern applications with unparalleled availability, protection, simplicity, and scale. It provides universal accessibility with support for object, and HDFS.

ECS Appliance enables cloud service providers to deliver competitive cloud storage services at scale. ECS provides a single platform for all web, mobile, Big Data, and social media applications.

upvoted 2 times

 **laureano21** 3 years, 9 months ago

Esta correcto

upvoted 2 times

 **haboliveira** 3 years, 6 months ago

Boas Laureano sera possivel entrar em contato ctg?

Humberto

upvoted 1 times

 **DragonTarzan** 4 years, 1 month ago

d:

<https://www.dell.com/community/s/vjauj58549/attachments/vjauj58549/solutions-ch/516/1/h13999-hadoop-ecs-data-services-wp.pdf>

upvoted 1 times



Which cloud deployment model allows provisioning of a cloud infrastructure only for use by a single organization having multiple consumers?

- A. Public cloud
- B. Multi-cloud
- C. Hybrid cloud
- D. Private cloud

Suggested Answer: *D*

Reference:

<https://www.paranet.com/blog/bid/128265/The-Four-Types-of-Cloud-Computing-Models>

  **oglcn** 3 years, 3 months ago

Private Cloud: Its a one-on-one environment for single user (Customer).

upvoted 1 times

A company plans to deploy a new application in their environment. The new application requires 2 TB of storage space. During peak workloads, the application is expected to generate 4500 IOPS with a typical I/O size of 4 KB. The capacity of each available disk drive is 250 GB. In addition, the maximum number of IOPS a drive can perform at with a 70% utilization is 90 IOPS.

What is the minimum number of disk drives needed to meet the application's capacity and performance requirements for a RAID 0 configuration?

- A. 8
- B. 40
- C. 50
- D. 70

Suggested Answer: D

Reference:

<https://www.dell.com/community/Student-Discussions/E10-001-Exam-Question-Query/td-p/7099780>

Community vote distribution

C (100%)


 **salmoallemi** Highly Voted 3 years, 9 months ago

IOPS needed is 4500

Each drive is 90 IOPS


So to reach the 4500 IOPS, we need 50 drive

upvoted 13 times

 **diaboliq1990** Highly Voted 3 years, 10 months ago

Answer is C

upvoted 9 times

 **th88780u980** Most Recent 1 year, 2 months ago

Ans : 50

To calculate the minimum number of disk drives needed for a RAID 0 configuration to meet the capacity and performance requirements, let's start by determining the number of drives required for capacity:

The application requires 2 TB of storage space, and each available disk drive has a capacity of 250 GB. So, to achieve 2 TB (2000 GB), you'd need $2000 \text{ GB} / 250 \text{ GB per drive} = 8$ disk drives.

Next, let's consider the IOPS requirement during peak workloads:

The application is expected to generate 4500 IOPS with a typical I/O size of 4 KB. Each drive can perform a maximum of 90 IOPS with a 70% utilization. To calculate the number of drives needed for IOPS, divide the required IOPS by the maximum IOPS per drive: $4500 \text{ IOPS} / 90 \text{ IOPS per drive} = 50$ drives

Therefore, for a RAID 0 configuration to meet both the capacity and performance requirements (considering the higher number between capacity and performance), you would need a minimum of 50 disk drives.

upvoted 1 times

 **NaelOsama** 2 years, 10 months ago

Selected Answer: C

For req performance (70% utilization - 90 IOPS): $4500/90=50$

upvoted 1 times

 **Ivan_Miran** 3 years, 4 months ago

Should be 50. IOPS are already defined at 70%.

"Consider an example in which the capacity requirement for an application is 1.46 TB. The number of IOPS generated by the application at peak workload is estimated at 9,000. IOPS. The vendor specifies that a 146-GB, 15,000-rpm drive is capable of doing a maximum 180 IOPS. In this example, the number of disks required to meet the capacity requirements will be $1.46 \text{ TB}/146 \text{ GB} = 10$ disks. To meet the application IOPS requirements, the number of disks required is $9,000/180 = 50$. <<However, if the application is response-time sensitive, the number of IOPS a disk drive can perform should be calculated based on 70-percent disk utilization. Considering this, the number of IOPS a disk can perform at 70

percent utilization is $180 \times 0.7 = 126$ IOPS. Therefore, the number of disks required to meet the application IOPS requirement will be $9,000/126 = 72.>>$

Information Storage and Management . Wiley. Kindle Editio.

upvoted 1 times

🗨️ **mozgachev** 3 years, 5 months ago

from course ISM v4:

Step 2: Calculate the maximum number of IOPS the drive can perform

- Maximum number of IOPS = $1 / 6.25 \text{ ms} = 160$ IOPS
- Maximum number of IOPS at 70% utilization = $160 \times 0.7 = 112$ IOPS

• Step 3: Calculate the number of drives for the application

- Drives required to meet performance requirement = $5200 / 112 = 47$
- Drives required to meet capacity requirement = $1.5 \text{ TB} / 250 \text{ GB} = 6$

• Number of drives required = Maximum (Capacity, Performance)

- Maximum (6, 47) = 47 disk drives

upvoted 1 times

🗨️ **Mikes81** 3 years, 7 months ago

Should be D , You want it to perform even with 100% utulization so add +25% slack

upvoted 5 times

🗨️ **runnerbean** 3 years, 3 months ago

@Mikes81 - The disks aren't running at 100% as the question says the IOPS per disk are worked out at 70%

upvoted 1 times

🗨️ **srichard** 3 years, 3 months ago

in this case, the utization mean 70% of disk capacity. it's the optimal rate for drives. at 100%, the disks performance decrease.

upvoted 1 times

Which technology decouples the OS, applications, and user state from a physical compute system?

- A. Application virtualization
- B. Storage virtualization
- C. Thin client virtualization
- D. Desktop virtualization

Suggested Answer: *D*

🗨️ 👤 **oglcn** 3 years, 3 months ago

Definition: Desktop Virtualization

Technology that decouples the OS, applications, and user state from a physical compute system to create a virtual desktop environment that can be accessed from any client device. (From ISMv4 Page:128)

upvoted 2 times

🗨️ 👤 **inf_ph** 3 years, 3 months ago

Hey oglcn,

Have you taken the exam?

upvoted 1 times

🗨️ 👤 **rustam_sh** 3 years, 5 months ago

correct

upvoted 1 times



What is the goal of a business continuity solution?

- A. Reduce the amount of investment in protecting an organization's business-related data.
- B. Avoid a disaster in a data protection environment to control the loss of data.
- C. Ensure information availability that is required to conduct vital business operations.
- D. Eliminate security threats in a data protection environment for continuous access to data.

Suggested Answer: C

Reference:

<https://www.oreilly.com/library/view/information-storage-and/9781118236963/9781118236963c09.xhtml>

  **gonglei** 1 year, 7 months ago

C

The goal of a BC solution is to ensure the "information availability" required to conduct vital business operation
upvoted 1 times

  **Expert2000** 3 years, 3 months ago

Business continuity (BC) is a set of processes that includes all activities that a business must perform to mitigate the impact of planned and unplanned downtime. BC entails preparing for, responding to, and recovering from a system outage that adversely affects business operations. It describes the processes and procedures an organization establishes to ensure that essential functions can continue during and after a disaster.
upvoted 2 times

What is a function of the metadata service in an object-based storage device (OSD)?

- A. Generates the file system namespace from the file contents.
- B. Manages both physical and logical disks on which the user data is stored.
- C. Manages the storage space for all remote clients.
- D. Maintains the object IDs and file system namespace mappings.

Suggested Answer: D

Community vote distribution

D (100%)

🗨️ 👤 **SupremeND** 2 years, 6 months ago

Selected Answer: D

The OSD node has two key services: metadata service and storage service. The metadata service is responsible for generating the object ID from the contents (may also include other attributes of data) of a file. It also maintains the mapping of the object IDs and the file system namespace. In some implementations, the metadata service runs inside an application server. The storage service manages a set of disks on which the user data is stored.

upvoted 1 times

🗨️ 👤 **Ivan_Miran** 3 years, 3 months ago

Answer D

The OSD system is composed of one or more nodes. A node is a server that runs the OSD operating environment and provides services to store, retrieve, and manage data in the system. The OSD node has two key services: metadata service and storage service. The metadata service is responsible for generating the object ID from the contents (and can also include other attributes of data) of a file.** It also maintains the mapping of the object IDs and the file system namespace. ** The storage service manages a set of disks on which the user data is stored. Information Storage and Management . Wiley. Kindle Edition

upvoted 1 times

DRAG DROP -

Match the types of attack with their respective description.

Select and Place:

Attack Types

Denial-of-Service
Phishing
Man-in-the-Middle
Hyperjacking

Description

Sends a spoofing email containing a link(s) to a fraudulent website
Installs a rogue hypervisor that takes control of a compute system
Prevents legitimate users from accessing resources or services
Eavesdrops on the network to capture credentials

Suggested Answer:

Attack Types

Denial-of-Service
Phishing
Man-in-the-Middle
Hyperjacking

Description

Phishing
Hyperjacking
Denial-of-Service
Man-in-the-Middle

 **kikigwee** 3 years, 1 month ago

Denial of service - prevents legitimate users from accessing resources or services

phising - send a spoofing email


Man in the middle - eavsdrops on the network


hyperjacking - install a rouge hypervisor that takes control of a computer system


upvoted 2 times

 **oglcn** 3 years, 3 months ago


Type of attack Description


Phishing  Social engineering attack used to deceive users


 Carried out by spoofing email containing link to a fake website

 Users credentials entered on the fake site are captured

Installing keystroke-logging malware

 Attacker installs malware in administrator's or user's compute system

 Malware captures users credentials and sends to the attacker

Man-in-the-middle  Attacker eavesdrops on the network to capture credentia

upvoted 1 times

  **AbhijatMindulkar** 3 years, 6 months ago

DoS Prevents legitimate users from accessing resources or services

Phishing Sends a spoofing email containing links to a fraudulent website

Man-in-the-middle Eavesdrops on the network to capture credentials

Hyperjacking Installs a rogue hypervisor that takes control of a compute system.

upvoted 3 times

What accurately describes virtual SAN (VSAN) trunking?

- A. Reduces the utilization of ISLs by removing all the ISLs configured between the FC switches.
- B. Increases the number of E_ports used for the ISLs to transfer multiple VSAN traffic.
- C. Increases the complexity of managing the FC SAN with an increase in the number of ISLs.
- D. Reduces the number of ISLs between switches that are configured with multiple VSANs.

Suggested Answer: *D*

Reference:

<https://www.mycloudwiki.com/san/fc-san-virtualization/>

 **Tyrion0097** 3 years, 6 months ago

Allows network traffic from multiple VSANs to traverse a single ISL (trunk link)

- Enables an E_Port (trunk port) to send or receive multiple VSAN traffic over a trunk link
 - Reduces the number of ISLs between switches that are configured with multiple VSANs
- upvoted 2 times

What does an area ID represent in Fibre Channel (FC) addressing?

- A. Group of switch ports
- B. Switch within a group
- C. Port within a group
- D. Group of switches


Suggested Answer: A


Community vote distribution

A (100%)

 **amtavt** Highly Voted 4 years, 5 months ago
Should be A

"The area ID is used to identify a group of switch ports that are used for connecting nodes."
upvoted 10 times

 **nicefella** 4 years, 3 months ago
any reference link?
upvoted 3 times

 **oglcn** 3 years, 3 months ago
You can find in ISMv4 page 318.
upvoted 1 times

 **Expert2000** Most Recent 3 years, 3 months ago
Selected Answer: A

A domain ID is a unique number that is provided to each switch in the fabric. The area ID is used to identify a group of switch ports that are used for connecting nodes.
upvoted 1 times


 **Ivan_Miran** 3 years, 4 months ago


"A domain ID is a unique number provided to each switch in the fabric. Although this is an 8-bit field, there are only 239 available addresses for domain ID because some addresses are deemed special and reserved for fabric management services. For example, FFFFFC is reserved for the name server, and FFFFFE is reserved for the fabric login service. <<The area ID is used to identify a group of switch ports >> used for connecting nodes. An example of a group of ports with a common area ID is a port card on the switch. The last field, the port ID, identifies the port within the group."

Information Storage and Management . Wiley. Kindle Edition.
upvoted 1 times

 **Saeid** 3 years, 9 months ago

A. a group of switch ports
Page 308: The area ID is used to identify a group of switch ports that are used for connecting nodes.
upvoted 2 times

 **AB1976** 3 years, 4 months ago
Can you get me a copy of the book? alwin bongers at hotmail com
upvoted 1 times

 **kens** 4 years, 1 month ago


The area ID is used to identify a group of F_ports. An example of a group of F_ports would be a card on the switch with more than one port on it
upvoted 1 times

 **atidd** 4 years, 5 months ago

For an N_port (Node Port), the first field in the address contains an 8-bit Domain ID field. Keep in mind, 239 of the possible Domain IDs are available for corporate use, while the rest are reserved. Following the Domain ID is an 8-bit Area ID, and then an 8-bit Port ID. The area ID is used to identify a group of F_ports (Fabric Ports), while the Port ID is used to identify the F_port within the group.

<https://www.cbttuggets.com/blog/technology/data/storage-mastery-fc-architecture>
upvoted 1 times

Based on the exhibit, what does the "X", "Y", and "Z" represent in a NAS environment?



- A. X = FC SAN, Y = IP Network, and Z = NAS Gateway
- B. X = IP Network, Y = FC SAN, and Z = NAS Gateway
- C. X = FC SAN, Y = NAS Gateway, and Z = IP Network
- D. X = IP Network, Y = NAS Gateway, and Z = FC SAN

Suggested Answer: B

Community vote distribution

B (100%)

andikovaci 3 years, 1 month ago

Selected Answer: B

correct

upvoted 1 times

[Removed] 3 years, 3 months ago

B is correct

upvoted 1 times

oglcn 3 years, 3 months ago

It should B. Correct.


upvoted 2 times

Ahmedalsharefy 3 years, 4 months ago

I think B or D

upvoted 1 times

Which type of zoning does the "X" represent?



- A. Fabric zoning
- B. WWN zoning
- C. Mixed zoning
- D. Port zoning

Suggested Answer: D

KZM 3 years, 7 months ago

Correct, Port zoning:

Port zoning uses the switch port ID to define zones.

upvoted 4 times

Abdiel 3 years, 8 months ago

Answer is correct = Port Zoning (Page 326)

Uses the switch port ID to define zones. In port zoning, access to node is determined by the physical switch port to which a node is connected. The zone members are the port identifiers (switch domain ID and port number) to which FC HBA and its targets (storage systems) are connected. If a node is moved to another switch port in the fabric, port zoning must be modified to enable the node, in its new port, to participate in its original zone. However, if an FC HBA or storage system port fails, an administrator has to replace the failed device without changing the zoning configuration.

upvoted 4 times

nate_slim 3 years, 10 months ago

Answer is correct = Port Zoning.

Port Zoning uses the switch port IDs to define the zones.

WWN Zoning uses the World wide Names to define the zones, this would show the zone stretch from the compute port to the storage port.

upvoted 3 times

DragonTarzan 4 years, 1 month ago

b:

What are the pros and cons of port zoning and WWN zoning on a Fibre Channel SAN fabric?

By Fibre Channel SAN zoning, we mean partitioning Fibre Channel fabric into groups to add security and improve management. Fibre Channel SANs can be zoned according to the World Wide Name (WWN) of each device, or according to switch ports.

WWN zoning groups a number of WWNs in a storage-area network zone and allows them to communicate with each other. The switch port that each device is connected to is irrelevant when WWN zoning is configured. An advantage to this type of zoning is that if a port is suspected of

being faulty, another port can be used without the need for fabric reconfiguration. A disadvantage with WWN zoning is that if there's a host bus adapter (HBA) failure, the fabric will need to be reconfigured for the host to reconnect to its storage. WWN zoning is sometimes referred to as soft zoning.

upvoted 1 times

Which functionality does a Control layer provide in a software-defined networking architecture?

- A. Defining the network behavior through various policies.
- B. Making decisions on how the packets should be forwarded.
- C. Establishing communications between the Infrastructure and Interface layers.
- D. Forwarding data packets based on the requirements from the Application layer.

Suggested Answer: B

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

I think it's B

"The architecture of SDN consists of three layers along with APIs in between to define the communication.

▮ Infrastructure Layer: This layer consists of networking devices such as switches and routers. It is responsible for handling data packets such as forwarding or dropping of packets and handling the devices. This layer forms the data plane and performs actions based on the instructions received.

▮ Control Layer: This layer consists of controllers and acts as a the brain of the SDN architecture. It is responsible for making decisions such as how the packets should be forwarded based on the requirements, and relays the decisions to the networking devices (data plane) for execution. It also extracts the information about the network from the data plane and communicates it to the application layer. This layer forms the control plane.

▮ Application Layer: This layer consists of applications and services such as business applications, and analytics that define the network behavior through policies and also define the requirements. It communicates the requirements through the APIs to the control layer. This layer forms the application plane of the SDN architecture."

upvoted 19 times

 **rachid1761960** Most Recent 3 years, 7 months ago

Application Layer: This layer consists of applications and services such as business applications, and analytics that define the network behavior through policies and also define the requirements. It communicates the requirements through the APIs to the control layer.

I think the correct is the D:

upvoted 2 times

 **techimoon4455** 4 years, 4 months ago

Lots of discussion on this. Great to see.

Any news about PowerStore exam questions /dumps?


DES-1221 Specialist Implementation Engineer PowerStore Solutions Exam

upvoted 1 times

 **[Removed]** 4 years, 4 months ago

well explained by amtavt - B

upvoted 3 times

 **learner** 4 years, 4 months ago

Agree B

upvoted 4 times

 **ameya84** 4 years, 5 months ago

C is correct

upvoted 1 times


What accurately describes an Incident in a storage infrastructure?

- A. Unknown root cause of one or more disruptions to service.
- B. Any event that may lead to the disruption of a service.
- C. Support request to help access an IT resource.
- D. Any event that may create a change in correcting a known error.

Suggested Answer: A

Community vote distribution

B (100%)

 **nemafa1615** Highly Voted 4 years, 5 months ago

Should be B

upvoted 17 times

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

"An incident is an unplanned event such as an HBA failure or an application error that may cause an interruption to services or degrade the service quality."

upvoted 8 times

 **SupremeND** Most Recent 2 years, 6 months ago

Selected Answer: B

B is the correct answer

upvoted 1 times

 **NaelOsama** 2 years, 10 months ago

Selected Answer: B

An incident is an unplanned event such as an HBA failure or an application error that may cause an interruption to services or degrade the service quality. Incident management is responsible for detecting and recording all incidents in a storage infrastructure. It investigates the incidents and provides appropriate solutions to resolve the incidents. It also documents the incident history with details of the incident symptoms, affected services, components and users, time to resolve the incident, severity of the incident, description of the error, and the incident resolution data. The incident history is used as an input for problem management (described next).

upvoted 2 times

 **andikovaci** 3 years, 1 month ago

Selected Answer: B

should be B

upvoted 1 times

 **beltaggy** 3 years, 4 months ago

B for me

upvoted 2 times

 **AbhijatMndulkar** 3 years, 6 months ago

B it is without a doubt.

upvoted 3 times

 **jeremy15n** 3 years, 7 months ago

"An incident is an UNPLANNED EVENT such as an HBA failure or an application error that may cause an interruption to services or degrade the service quality."

It should be A because B is ANY event and A is an UNKNOWN event which best describes the definition from the ISMv4 guide.

upvoted 2 times

 **Mikes81** 3 years, 7 months ago

Power outage in a data center is not "unkown" but a huge incident so A is wrong.

Should be B

upvoted 2 times

🗨️ 👤 **shaolong1231** 3 years, 10 months ago

Can you share the ISMv4 guide. Thanks
upvoted 2 times

🗨️ 👤 **visgoy** 3 years, 12 months ago

should be b
upvoted 4 times

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

page 785 of the ISMv4 student guide

An incident is an unplanned event such as an HBA failure or an application error that may cause an interruption to services or degrade the service quality.

upvoted 4 times

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

Please share the student guide with me
upvoted 3 times

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

Can you share the ISMv4 guide. Thanks
upvoted 2 times

What accurately describes the scale-out storage architecture?

- A. Scales the capacity only within a single storage system.
- B. Provides a fixed capacity ceiling that limits scalability beyond a storage system.
- C. Provides the capability to maximize its capacity by adding nodes to the cluster.
- D. Restricts scalability to the limits of the storage controllers.

Suggested Answer: *C*

Reference:

<https://whatis.techtarget.com/definition/scale-out-storage>

🗨️ 👤 **Saroj1969** 2 years, 8 months ago

correct

upvoted 1 times

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

C correct

upvoted 4 times

🗨️ 👤 **rustam_sh** 3 years, 5 months ago

correct

upvoted 2 times

Which technology combines NAND flash memory into a hard disk drive and is a single, integrated device incorporating both types of storage?

- A. Solid-state drive
- B. Dual drive
- C. Fibre Channel drive
- D. Solid-state hybrid drive

Suggested Answer: *D*

Reference:

https://en.wikipedia.org/wiki/Hybrid_drive#Solid-state_hybrid_drive

🗨️ 👤 **Saroj1969** 2 years, 8 months ago

correct

upvoted 1 times

🗨️ 👤 **rustam_sh** 3 years, 5 months ago

correct

upvoted 1 times

Which Dell EMC product creates a server and IP-based SAN from direct-attached server storage to deliver flexible and scalable performance and capacity on demand?

- A. PowerMax
- B. VxFlex OS
- C. ViPR Controller
- D. XtremIO

Suggested Answer: B

Reference:

https://www.dell EMC.com/mk-mk/collaterals/unauth/white-papers/products/ready-solutions/h17332_dell EMC_vxflex_os_networking_best_practices.pdf

  **Saroj1969** 2 years, 8 months ago

correct

upvoted 1 times

  **Quntana70** 3 years, 1 month ago

Dell EMC VxFlex OS

Software that creates a server and IP-based SAN from direct-attached server storage to deliver flexible and scalable performance and capacity on demand. As an alternative to a traditional SAN infrastructure, VxFlex OS combines HDDs, SSDs, and PCIe flash cards to create a virtual pool of block storage with varying performance tiers. It decouples compute and storage, and scales each resource together or independently to drive maximum efficiency and to eliminate wasted CAPEX at scale. Distributed I/O Parallelism vs. Data Locality: uses all resources to deliver against all I/O requests to drive massive performance. Eliminates bottlenecks and scales performance linearly.

upvoted 1 times

  **rustam_sh** 3 years, 5 months ago

correct

upvoted 2 times

Which Internet of Things (IoT) component is responsible for the automation of operations by applying a force based on the dynamics of data?

- A. Actuator
- B. Switch
- C. Sensor
- D. Gateway

Suggested Answer: C

Community vote distribution

A (100%)

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

It's A. Actuator :

"Actuators: Devices that collect data from sensors and perform the required action. Actuators consume energy to produce physical action like creating a motion or controlling a system. Examples of actuators include electric motor that uses electric power to generate motion, hydraulic actuators use fluid pressure to generate motion. In IoT, actuators help to automate the operations by applying a force based on the dynamics of data generated by sensors."

upvoted 16 times

  **van2020** Highly Voted 4 years, 4 months ago

it should be A.

Actuators help to automate the operations by applying a force based on the dynamics of data generated by sensors.

Sensors should be able to detect a wide range of physical phenomenon ranging from temperature, pressure, to motion and magnetic fields.



upvoted 8 times

  **anirbansarkar1069** Most Recent 2 years, 10 months ago

Selected Answer: A

Should be A.

upvoted 2 times

  **FenrirIII** 3 years, 2 months ago

Selected Answer: A

Should be an Actuator

upvoted 2 times

  **Ivan_Miran** 3 years, 4 months ago

Although correct answer should be "A", on the practice test (DELL Academy) they answered it as sensor.

upvoted 1 times

  **AbhijatMndulkar** 3 years, 5 months ago



The answer is A. actuator and not C

upvoted 3 times

  **luchiano33** 3 years, 11 months ago

Its A Actuator

upvoted 1 times

  **visgoy** 3 years, 12 months ago

should be a

upvoted 1 times

In an FCoE switch, which function is provided by the Fibre Channel Forwarder (FCF)?

- A. Communication bridge between CEE and IP networks
- B. Communication bridge between FC and FCIP networks
- C. Communication bridge between CEE and FCIP networks
- D. Communication bridge between CEE and FC networks



Suggested Answer: D

- 🗨️ **Domina** Highly Voted 4 years, 5 months ago
FCF functions as the communication bridge between CEE and FC networks
upvoted 14 times
- 🗨️ **ssajith** 4 years, 1 month ago
FCF functions as the communication bridge between CEE and FC networks\
upvoted 1 times
- 🗨️ **ameya84** Highly Voted 4 years, 5 months ago
D is correct
upvoted 11 times
- 🗨️ **Saroj1969** Most Recent 2 years, 8 months ago
it is D
upvoted 1 times
- 🗨️ **Ivan_Miran** 3 years, 3 months ago
Answer D:
An FCoE switch has both Ethernet switch and Fibre Channel switch functionalities. The FCoE switch has a Fibre Channel Forwarder (FCF), Ethernet Bridge, and set of Ethernet ports and optional FC ports, as shown in Figure 6.15. The function of the FCF is to encapsulate the FC frames, received from the FC port, into the FCoE frames and also to de-encapsulate the FCoE frames, received from the Ethernet Bridge, to the FC frames
upvoted 1 times
- 🗨️ **JoZi_Berlin** 3 years, 7 months ago
It is B I think.
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/8_x/config/fcoe/cisco_mds9000_fcoe_config_guide_8x/overview.pdf

There is FCoE in the question, not CEE.
upvoted 1 times
- 🗨️ **JoZi_Berlin** 3 years, 7 months ago
Sorry, this was wrang. 😊
As stated below, the answer can be found on page 401 in the ISMv4 guide. It is D.
upvoted 1 times
- 🗨️ **Peach** 4 years, 1 month ago
FCF functions as the communication bridge between CEE and FC networks
It is D
upvoted 2 times
- 🗨️ **MSPB** 4 years, 1 month ago
Page 401 ISM V4.0 manual (Student Guide)
upvoted 4 times
- 🗨️ **Flintk** 4 years, 4 months ago
🗨️
FCF functions as the communication bridge between CEE and FC networks
🗨️ Encapsulates and decapsulates FC frames

C is Correct

upvoted 1 times

  **EW3772** 4 years, 3 months ago

if FC networks, it should be "D"

upvoted 3 times

Which technology is used in the media and entertainment industry to enable content-based search activity and to create video subtitles using natural language processing?



- A. Machine learning
- B. Social networking
- C. Cloud computing
- D. Internet of Things

Suggested Answer: A

  **Saroj1969** 2 years, 8 months ago

A is correct



upvoted 1 times

  **moskacgh** 3 years, 8 months ago

On page 91 of ISMv4 student guide

Content from the media and entertainment industry can be automatically tagged using metadata by applying machine learning solutions.

upvoted 2 times

  **Kate321** 3 years, 11 months ago

Content from the media and entertainment industry can be automatically tagged using metadata by applying machine learning solutions. This method enhances content-based search activity by finding the right content quickly and helps the content developers to optimize the content to specific audiences based on their search data.

upvoted 1 times

Why are organizations moving towards a modern data center implementation?

- A. Develop modern applications in the public cloud without having the right skill set in-house.
- B. Ability to be agile, operate in real time, and develop intelligent products.
- C. Deploy only unified and converged IT resources to reduce the management complexity.
- D. Reduce the cost of managing data center components by using a self-service catalog.

Suggested Answer: C

Reference:

<https://www.actualtechmedia.com/wp-content/uploads/2018/05/Building-a-Modern-Data-Center-ebook.pdf>

Community vote distribution


B (100%)

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

It's B

Digital transformation is disrupting every industry, and with the evolution of modern technologies, organizations are facing too many business challenges. Organizations must operate in real time, develop smarter products, and deliver a great user experience. They must be agile, operate efficiently, and make decisions quickly to be successful.


upvoted 36 times

 **SupremeND** Most Recent 2 years, 6 months ago

Selected Answer: B

Digital transformation is disrupting every industry, and with the evolution of modern technologies, organizations are facing too many business challenges. Organizations must operate in real time, develop smarter products, and deliver a great user experience. They must be agile, operate efficiently, and make decisions quickly to be successful. However, these disruptive technologies along with agile methodologies are less resilient on traditional IT infrastructure and services. Organization's IT department also faces several challenges in supporting business challenges. So, organizations are moving towards modern data center to overcome the business challenges and be successful in their digital transformation journey.

upvoted 1 times

 **KimZhou** 2 years, 10 months ago

Selected Answer: B

Digital transformation is disrupting every industry, and with the evolution of modern technologies, organizations are facing too many business challenges. Organizations must operate in real time, develop smarter products, and deliver a great user experience. They must be agile, operate efficiently, and make decisions quickly to be successful.

upvoted 1 times

 **Nourhane** 3 years, 2 months ago

Organizations are moving towards modern data center to overcome the business and IT challenges (ISMv4 page 41)

upvoted 1 times

What is a key objective to performing continuous replication?

- A. Avoid performing data archiving.
- B. Reduce the RPO to near-zero.
- C. Reduce the deduplication ratio.
- D. Avoid performing data backups.

Suggested Answer: B

Community vote distribution

B (100%)

ameya84 **Highly Voted** 4 years, 5 months ago

B is correct

upvoted 15 times

van2020 **Highly Voted** 4 years, 4 months ago

Yeah it should be B.

The data on the replica is in-sync with the production data always. The objective with any continuous replication is to reduce the RPO to zero or near-zero.

upvoted 5 times

SupremeND **Most Recent** 2 years, 6 months ago

Selected Answer: B

The data on the replica is in-sync with the production data always. The objective with any continuous replication is to reduce the RPO to zero or near-zero.

upvoted 1 times

Saroj1969 2 years, 8 months ago

B is correct

upvoted 1 times

jebac 4 years, 2 months ago

B is right

upvoted 5 times

robinMcAfee 4 years, 2 months ago

B is right

upvoted 3 times

Jamestyler22 4 years, 4 months ago

The data on the replica is in-sync with the production data always. The objective with any continuous replication is to reduce the RPO to zero or near-zero.

B is correct

upvoted 2 times

What accurately describes a shared pool in virtual provisioning?

- A. Should contain only homogeneous disk drive types.
- B. Supports a single RAID protection level.
- C. Consists of only a virtual storage volume.
- D. Should contain only physical drives with different sizes.

Suggested Answer: B

🗨️ **Ivan_Miran** 3 years, 4 months ago

B

"Physical storage is allocated to the host "on-demand" from a shared pool of physical capacity. A shared pool consists of physical disks. A shared pool in virtual provisioning is analogous to a RAID group, which is a collection of drives on which LUNs are created. <<Similar to a RAID group, a shared pool supports a single RAID protection level>>. However, unlike a RAID group, a shared pool might contain large numbers of drives. Shared pools can be homogeneous (containing a single drive type) or heterogeneous (containing mixed drive types, such as flash, FC, SAS, and SATA drives).

Information Storage and Management . Wiley. Kindle Edition.

upvoted 2 times

🗨️ **moskacgh** 3 years, 8 months ago

Correct: B

A shared pool in virtual provisioning is analogous to a RAID set, which is a collection of drives on which LUNs are created. Similar to a RAID set, a shared pool supports a single RAID protection level.

upvoted 2 times

🗨️ **atidd** 4 years, 5 months ago

B is correct.

Shared pools can be homogeneous or heterogeneous when it comes to drive types. But single RAID protection level.

upvoted 3 times

Which data archiving component scans primary storage to find the files that meet the archiving policy?

- A. Archive storage
- B. Archive database server
- C. Archive agent
- D. Archive stub file

Suggested Answer: A

Community vote distribution

C (100%)

🗨️ **Seahawk270** Highly Voted 4 years, 5 months ago

SHould be C

"Archiving agent scans primary storage to find files that meet the archiving policy. The archive server indexes the files."

upvoted 25 times

🗨️ **MSPB** Highly Voted 4 years, 1 month ago

C see student guide page 639

upvoted 7 times

🗨️ **Saroj1969** Most Recent 2 years, 8 months ago

c agent

upvoted 1 times

🗨️ **Quntana70** 3 years, 1 month ago

C The data archiving operation has an archiving agent, archive server/policy engine, and archive storage. The archiving agent scans the primary storage to find files that meet the archiving policy. This policy is defined on the archive server (policy engine).

upvoted 1 times

🗨️ **oumima** 3 years, 1 month ago

Archiving agent scans primary storage to find files that meet the archiving policy. The archive server indexes the files.

upvoted 1 times

🗨️ **FenrirIII** 3 years, 2 months ago

Selected Answer: C

Should be C: The archiving agent scans primary storage to find files that meet the archiving policy.

upvoted 1 times

🗨️ **bzielazny** 3 years, 2 months ago

Selected Answer: C

Archiving agent scans primary storage to find files that meet the archiving policy. The archive server indexes the files

upvoted 1 times

🗨️ **Ivan_Miran** 3 years, 3 months ago

Answer C

Archiving solution architecture consists of three key components: archiving agent, archiving server, and archiving storage device. <<An archiving agent is software installed on the application server. The agent is responsible for scanning the data that can be archived based on the policy defined on the archiving server. >> After the data is identified for archiving, the agent sends the data to the archiving server. Then the original data on the application server is replaced with a stub file. The stub file contains the address of the archived data. The size of this file is small and significantly saves space on primary storage. This stub file is used to retrieve the file from the archive storage device.

Information Storage and Management . Wiley. Kindle Edition.

upvoted 1 times

🗨️ 👤 **visgoy** 3 years, 12 months ago

should be C

upvoted 2 times

🗨️ 👤 **Peach** 4 years, 1 month ago

Archiving agent scans primary storage to find files that meet the archiving policy. The archive server indexes the files.

So it is C

upvoted 3 times

In the event of a disk failure, which RAID level enables quick data reconstruction?

- A. RAID 3
- B. RAID 5
- C. RAID 6
- D. RAID 10

Suggested Answer: C

Community vote distribution

D (75%)

B (25%)

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

I think it should be D
upvoted 17 times

 **mozgachev** 3 years, 5 months ago

yes, D

RAID 5 and 6 use algorithms for a rebuild. RAID 5 use XOR this is fast because 1 math action for 1 block rebuild, RAID 6 use XOR and algorithm Reed Solomon is slowly because 2 math actions for rebuild 1 block data. The RAID 1 has a mirror disk then rebuild is only copy back.

upvoted 2 times

 **[Removed]** Highly Voted 4 years, 4 months ago

D or RAID10 has the fastest rebuild time.

<https://www.prepressure.com/library/technology/raid>


upvoted 9 times

 **JRO1** Most Recent 2 years, 10 months ago

Selected Answer: D

All RAID levels have to "reconstruct" once the failed drive is replaced to re-establish resiliency. A RAID 10 does not have to calculate data using any parity....it simply copies the data to a mirror.

upvoted 1 times

 **superbeq** 2 years, 10 months ago

Selected Answer: D

Answer is D

upvoted 1 times

 **NaelOsama** 2 years, 10 months ago

Selected Answer: B

The rebuild operation in RAID 6 may take longer than that in RAID 5 due to the presence of two parity sets

upvoted 1 times

 **andikovaci** 3 years, 1 month ago

Selected Answer: D

<https://www.deft.com/blog/the-levels-of-raid/>

upvoted 1 times

 **Quntana70** 3 years, 1 month ago

Should be 10

upvoted 1 times

 **Insua** 3 years, 5 months ago

The question correct is RADI 5 The rebuild operation in RAID 6 may take longer than that in RAID 5 due to the presence of two parity sets.

upvoted 1 times

 **Gabeeee** 3 years, 7 months ago

It's gonna be D
upvoted 1 times

🗨️ 👤 **Mahrizou** 3 years, 10 months ago
it couldn't be RAID6:

Notes

RAID 6 distributes the parity across all the disks. The write penalty (explained later in this module) in RAID 6 is more than that in RAID 5; therefore, RAID 5 writes perform better than RAID 6.

The rebuild operation in RAID 6 may take longer than that in RAID 5 due to the presence of two parity sets.

upvoted 2 times

🗨️ 👤 **visgoy** 3 years, 12 months ago
lingo is confusing - I will go with RAID 10
upvoted 1 times

🗨️ 👤 **JBONINO** 4 years ago
Advantages of RAID 10

If something goes wrong with one of the disks in a RAID 10 configuration, the rebuild time is very fast since all that is needed is copying all the data from the surviving mirror to a new drive. Should be D

upvoted 3 times

🗨️ 👤 **DragonTarzan** 4 years, 1 month ago
RAID10 has the fastest rebuild time
upvoted 2 times

🗨️ 👤 **Mark27591** 4 years, 1 month ago
Page 227 of the exam guide says: The rebuild operation in RAID 6 may take longer than that in RAID 5 due to the presence of two parity sets. So could be RAID 5, but probably 10.
upvoted 1 times

🗨️ 👤 **jackhuang** 4 years, 4 months ago
Raid6 should be accurate.
upvoted 1 times

🗨️ 👤 **Gruchmen** 4 years, 5 months ago
Probably not D because with RAID10 there is no need to reconstruct data - data is in complete, we only loose redundant data. With RAID5/6 there is no redundant data and lost data has to be re-calculated from the parity info.
upvoted 4 times

What is an advantage of implementing file-level virtualization in a data center environment?

- A. Provides non-disruptive file mobility across file servers or NAS devices.
- B. Provides automated scaling of file systems across NAS servers.
- C. Provides the capability to access files from any device.
- D. Provides dynamic distribution of files across NAS clients.

Suggested Answer: A

  **moskacgh** Highly Voted 3 years, 8 months ago

Correct: A

File-level virtualization, which is implemented in NAS or the file server environment, provides a simple, non-disruptive file-mobility solution.

upvoted 6 times

  **Ivan_Miran** Most Recent 3 years, 3 months ago

Answer A

"File-level virtualization, implemented in the file sharing environment, provides a simple, nondisruptive file-mobility solution. It enables the movement of files across NAS devices, even if the files are being accessed.

...

File-level virtualization simplifies file mobility. It provides user or application independence from the location where the files are stored.

Information Storage and Management . Wiley. Kindle Edition.



upvoted 3 times

  **inf_ph** 3 years, 3 months ago

Hey Ivan_Miran

Did you take the exam?

upvoted 2 times

  **Kate321** 3 years, 11 months ago

A is correct File-level virtualization

facilitates the movement of files across online file servers or NAS devices.

upvoted 3 times

DRAG DROP -





To authenticate a user using OpenID authentication control, what are the correct sequence of steps?

Select and Place:

User provides consent to profile sharing	STEP 1
User sends a login request using OpenID	STEP 2
Authentication response is redirected to the organization providing the services	STEP 3
Authentication request is redirected to an OpenID provider	STEP 4

Suggested Answer:

User provides consent to profile sharing	User sends a login request using OpenID
User sends a login request using OpenID	Authentication request is redirected to an OpenID provider
Authentication response is redirected to the organization providing the services	Authentication response is redirected to the organization providing the services
Authentication request is redirected to an OpenID provider	User provides consent to profile sharing

-  **Kimtetsu** Highly Voted 4 years ago
 Step 1: Login request using OpenID
 Step 2: Authentication request is redirected to OpenID provider
 Step 3: Consent to profile sharing
 Step 4: Authentication response is redirected to organization providing services
 upvoted 28 times
-  **Shadyamy** Highly Voted 3 years, 11 months ago
 Step 1: Login request using OpenID
 Step 2: Authentication request is redirected to OpenID provider
 Step 3: Consent to profile sharing
 Step 4: Authentication response is redirected to organization providing services
 upvoted 5 times
-  **Maryjn3** Most Recent 3 years, 1 month ago
 USER SENDS A LOGIN REQUEST USING OPEN ID
 AUTHENTICATION REQUEST IS REDIRECTED TO AN OPENID PROVIDER
 USER PROVIDES CONSENT TO PROFILE SHARING
 AUTHENTICATION REQUEST IS REDIRECTED TO AN OPENID PROVIDER
 upvoted 1 times
-  **David_Le** 3 years, 2 months ago
 Step 1: User sends a login request using OpenID
 Step 2: Authentication request is redirected to OpenID provider
 Step 3: User provides consent to profile sharing
 Step 4: Authentication response is redirected to organization providing services
 upvoted 3 times

Which product delivers virtualized or hosted desktops and applications through a single platform?

- A. VMware Horizon
- B. Dell EMC VxRail Appliance
- C. VMware NSX
- D. Dell PowerEdge servers

Suggested Answer: A

Reference:

<https://www.dell.com/pk/business/p/vmware-horizon-7/pd>

 **JBONINO** Highly Voted 4 years ago

VMware Horizon is a VDI solution for delivering virtualized or hosted desktops and applications through a single platform to the end users. These desktop and application services—including RDS, hosted apps, packaged apps with VMware ThinApp, and SaaS apps—can all be accessed from one unified workspace across devices and locations. Horizon provides IT with a streamlined approach to deliver, protect, and manage desktops and applications while containing costs and ensuring that end users can work anytime, anywhere, on any device. Horizon supports both Windows as well as Linux-based desktops.

upvoted 8 times

 **Expert2000** Most Recent 3 years, 3 months ago

VMware Horizon = Desktop Virtualization Product From VMware

upvoted 2 times

Which product provides an organization with the power to search and analyze vast amounts of data in one place and helps administrators to discover suspicious event trends?

- A. Dell InTrust
- B. Dell Change Auditor
- C. VMware AppDefense
- D. VMware AirWatch

Suggested Answer: A

Reference:

<https://cybersecurity-excellence-awards.com/candidates/dell-intrust/>

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

Dell Change Auditor

- Helps customer audit, alert, protect, and reports user activity and configuration
- The software has role-based access
- Enables customers to see how data is being handled

Dell InTrust

- Provides the organizations the power to search and analyze vast amounts of data in one place
- Provides information on who accessed the data, how was it obtained and how the data was used

VMware Airwatch

- Enables secure access to corporate resources
- Configures and updates device settings over-the-air, and secures mobile devices
- Manages different types of devices from a single console

VMware AppDefense

- Provides data center endpoint security
- Supports integration with third parties
- Provides automatic response
- Secures modern application

upvoted 15 times

 **David_Le** Most Recent 3 years, 2 months ago

A. Dell InTrust

upvoted 2 times

A company's Finance department needs 200 GB of storage for their business applications. A LUN is created with 200 GB and RAID 1 is configured to support the applications. If the cost of storage per GB is \$2, what is the chargeback cost to the department for application storage?

- A. \$ 200
- B. \$ 400
- C. \$ 800
- D. \$ 1500


Suggested Answer: C

 **JBONINO** Highly Voted 4 years ago

200GB a 2\$ el GB = 400\$ pero raid 1 utiliza doble del espacio => 400GB a 2\$ =800\$
upvoted 8 times

 **David_Le** Most Recent 3 years, 2 months ago

The needed raw capacity is 400GB, so the cost will be \$800.
upvoted 1 times

 **KZM** 3 years, 7 months ago

For RAID 1, 400 GB Raw capacity is needed for usable capacity 200 GB. So total cost of the application > 400*\$2 = \$ 800
upvoted 4 times

What is an example of a Performance Management task?

- A. Selecting an appropriate RAID type for a storage pool.
- B. Removing unused resources from a service and reassigning them to another.
- C. Scanning applications and databases to identify vulnerabilities.
- D. Planning and architecting data backup and replication solutions.

Suggested Answer: B


Community vote distribution

A (100%)

 **amtavt** Highly Voted 4 years, 5 months ago
It's A

'The storage system-related tasks include separating sequential and random I/Os to different spindles, selecting an appropriate RAID type for a storage pool, and changing storage tiering policy and cache configuration, when the performance management is concerned.'

upvoted 18 times

 **KingPo** Most Recent 2 years, 8 months ago

Selected Answer: A

B) Removing unused resources... is by ISM official document "Capacity Management"
A) Selecting RAID Level is shown for Performance Management
upvoted 1 times

 **kikigwee** 3 years, 1 month ago

Selected Answer: A

Key functions:

- Measures and analyzes the response time and throughput of components
- Identifies components that are performing below the expected level
- Makes configuration changes to optimize performance and address issues

Examples of performance management activities:

- Tuning database design, resource allocation to VMs, and multipathing
- Adding new ISLs and aggregating links to eliminate bottleneck
- Separating sequential and random I/Os to different spindles
- Changing storage tiering policy and cache configuration

upvoted 2 times

 **andikovaci** 3 years, 1 month ago

Selected Answer: A

a is correct

upvoted 1 times


 **David_Le** 3 years, 2 months ago

A. Selecting an appropriate RAID type for a storage pool.
upvoted 1 times

 **DragonTarzan** 4 years, 1 month ago

A it is obvious


upvoted 4 times

 **Peach** 4 years, 1 month ago

The performance management tasks on a SAN include implementing new ISLs and aggregating links in a multiswitch fabric to eliminate performance bottleneck. The storage system-related tasks include separating sequential and random I/Os to different spindles, selecting an appropriate RAID type for a storage pool, and changing storage tiering policy and cache configuration, when the performance management is concerned.

Answer A is confirmed

upvoted 3 times

 **Varunraj** 4 years, 3 months ago

Yes. Its A

upvoted 2 times

DRAG DROP -

If the client requests authorization from the resource owner, what are the remaining sequence of steps to authenticate a user using OAuth control?

Select and Place:

Client requests an access token by authenticating with the authorization server and presents the authorization grant

Client requests the protected resource from the resource server and authenticates by presenting the access token

Client receives an authorization grant used by the client to obtain an access token

Resource server validates the access token and, if valid, serves the request

Authorization server authenticates the client and validates the authorization grant

STEP 1

STEP 2

STEP 3

STEP 4

STEP 5

Suggested Answer:

Client requests an access token by authenticating with the authorization server and presents the authorization grant

Client requests the protected resource from the resource server and authenticates by presenting the access token

Client receives an authorization grant used by the client to obtain an access token

Resource server validates the access token and, if valid, serves the request

Authorization server authenticates the client and validates the authorization grant

Client requests the protected resource from the resource server and authenticates by presenting the access token

Client receives an authorization grant used by the client to obtain an access token


Client requests an access token by authenticating with the authorization server and presents the authorization grant

Authorization server authenticates the client and validates the authorization grant

Resource server validates the access token and, if valid, serves the request

Reference:

<https://tools.ietf.org/html/rfc6749>

 **shaolong1231** Highly Voted 3 years, 10 months ago

correct :

STEP1 : Client receives an authorization grand used by the client to obtain an access token

STEP2 : client requests an access token by authenticating with the authorization server and presenting the authorization grant

STEP3 : authorization server authenticates the client and validates the authorization grant

STEP4 : client requests the protected resource from the resource server and authenticates by presenting the access token

STEP5 : resource server validates the access token and if valid serves the request

<https://tools.ietf.org/html/rfc6749#section-7> P7

upvoted 10 times

 **marmaris** Highly Voted 4 years ago

Wrong

Client requests an authorization grand used by the client to obtain an access token

Client requests an access token by authenticating with the authorization server and presents the auth. Grand

Authorization server authenticates the client and validates the authorization grant

Client requests the protected resource from the res. Server and authenticates by pres. The access token.

6

Resource server validates the access token and serves the request.

upvoted 8 times

  **Saroj1969** Most Recent 2 years, 8 months ago

ans is not correct

upvoted 1 times

  **Maryjn3** 3 years, 1 month ago

CLIENT REQUESTS AN ACCESS TOKEN BY AUTHENTICATING WITH THE AUTHORIZATION SERVER AND PRESENTS THE AUTHORIZATION GRANT

CLIENT RECEIVES AN AUTHORIZATION GRANT USED BY THE CLIENT TO OBTAIN AN ACCESS TOKEN

AUTHORIZATION SERVER AUTHENTICATES THE CLIENT AND VALIDATES THE AUTHORIZATION GRANT

CLIENT REQUESTS THE PROTECTED RESOURCE FROM THEIR RESOURCE SERVER AND AUTHENTICATES BY PRESENTING THE ACCESS TOKEN

RESOURCE SERVER VALIDATES THE ACCESS TOKEN AND, IF VALID, SERVES THE REQUEST

upvoted 1 times

  **David_Le** 3 years, 2 months ago

STEP 1: client receives an authorization grand used by the client to obtain an access token

STEP 2: client requests an access token by authenticating with the authorization server and presenting the authorization grant

STEP 3: authorization server authenticates the client and validates the authorization grant

STEP 4: client requests the protected resource from the resource server and authenticates by presenting the access token

STEP 5: resource server validates the access token and if valid serves the request

upvoted 1 times

  **dioamm** 3 years, 3 months ago

ISMv4:

1. The client requests authorization from the resource owner. The authorization request can be made directly to the resource owner, or indirectly through the authorization server.

2. The client receives an authorization grant, which is a credential representing the resource owner's authorization to access its protected resources. It is used by the client to obtain an access token. Access tokens are credentials that are used to access protected resources. An access token is a string representing an authorization issued to the client. The string is usually opaque to the client. Tokens represent specific scopes and durations of access, granted by the resource owner, and enforced by the resource server and authorization server.

3. The client requests an access token by authenticating with the authorization server and presenting the authorization grant.

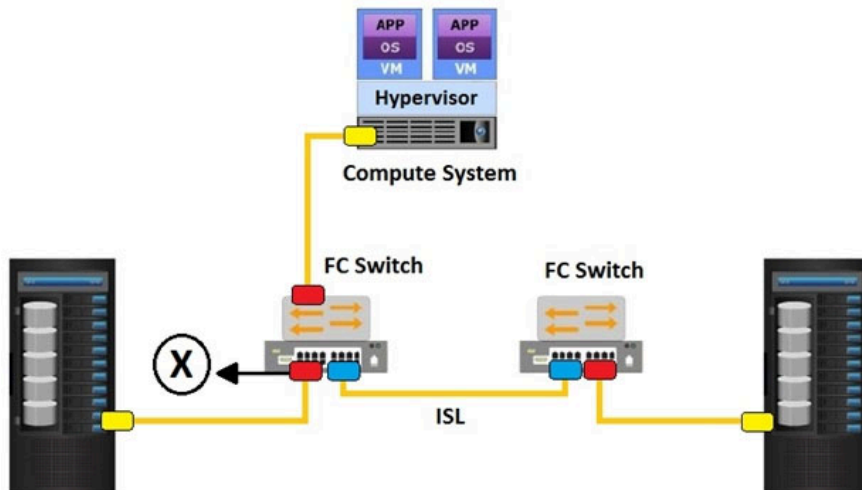
4. The authorization server authenticates the client and validates the authorization grant, and if valid, issues an access token.

5. The client requests the protected resource from the resource server and authenticates by presenting the access token.

6. The resource server validates the access token, and if valid, serves the request.

upvoted 2 times

What does the "X" represent in the exhibit?



- A. E_port
- B. F_port
- C. N_port
- D. NL_port

Suggested Answer: C

Community vote distribution

B (100%)

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

Selected Answer: B

F port
upvoted 1 times

🗨️ **Ersin** 3 years ago

F_port: A port on a switch that connects an N_Port. It is also known as a fabric port

Answer is B
upvoted 1 times

🗨️ **David_Le** 3 years, 2 months ago

It should be F_port.
upvoted 1 times

🗨️ **Ivan_Miran** 3 years, 3 months ago

Should be B:

"Ports in a switched fabric can be one of the following types:

N_Port: An end point in the fabric. This port is also known as the node port. Typically, it is a host port (HBA) or a storage array port connected to a switch in a switched fabric.

E_Port: A port that forms the connection between two FC switches. This port is also known as the expansion port. The E_Port on an FC switch connects to the E_Port of another FC switch in the fabric through ISLs.

F_Port: A port on a switch that connects an N_Port. It is also known as a fabric port.

G_Port: A generic port on a switch that can operate as an E_Port or an F_Port and determines its functionality automatically during initialization.

Information Storage and Management . Wiley. Kindle Edition.

upvoted 4 times

🗨️ **Drakker** 3 years, 5 months ago

B. F_port ISM page 323

upvoted 3 times

🗨️ **manny24** 3 years, 5 months ago

Its F_Port - that's the right answer

upvoted 3 times

🗨️ **dannyrad** 3 years, 5 months ago

In the diagram,

- 1) F_Ports are marked in Red. Remember, all ports at the FC Switch/Fabric are F ports.
- 2) N_Ports are marked in Yellow. The ones going towards F_Ports on the client/servers are N_Ports
- 3) E_Ports are marked in yellow. Extension Ports are switch-to-switch.

upvoted 3 times

🗨️ **andikovaci** 3 years, 1 month ago

I think you mean

E_Ports are market in blue. the port that connect between 2 FC switches.

upvoted 1 times

🗨️ **Bastiaan** 3 years, 9 months ago

It's F because they are referring to the switch port.

upvoted 3 times

🗨️ **salmoallemi** 3 years, 9 months ago

It is N_Port

F_port is switch port connects to N_Port. The answer is correct as it is.

upvoted 2 times

🗨️ **r8derfan33** 3 years ago

but this port is on the switch...should be F port.

upvoted 1 times

🗨️ **salmoallemi** 3 years, 9 months ago

C is correct

upvoted 2 times

🗨️ **suzcz** 3 years, 10 months ago

F_port is the ans

upvoted 3 times

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

F_Port = Fabric port ("source" port that connects to an N_Port)

N_Port = Node port (Destination port, e.g. the storage System or the hypervisor)

E_Port = Expansion port (Connect to another switch)

G_Port = Generic port (Can be used as either F or E)

upvoted 2 times

🗨️ **Shadyamy** 3 years, 11 months ago

Please change answer to B

upvoted 1 times

🗨️ **visgoy** 3 years, 12 months ago

f port

upvoted 3 times

🗨️ **JBONINO** 4 years ago

F_port (RED on Exhibit). A port on a switch that connects an N_Port. It is also known as a fabric port.

upvoted 4 times

🗨️ **io_58** 4 years ago

shouldn't the answer be F_port? they are asking the name of port in fabric switch and F_port in switch connects to N_port in compute system?

upvoted 1 times

🗨️ **Ddd45** 4 years, 1 month ago



It is F_Port check page 229 figure

upvoted 2 times

  **renzitonoda13** 4 years, 1 month ago

Sure, what is N_port or F_port??

upvoted 1 times

  **eddb** 4 years, 1 month ago

F_port

upvoted 1 times

A backup administrator has defined a backup policy. The policy requires full backups to be performed at 10 PM every Sunday and incremental backups performed at 10 PM the remaining days. If files were accidentally deleted Wednesday morning, how many backup copies are required to restore the files?

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

Suggested Answer: A

Community vote distribution

B (100%)

🗨️ **Seahawk270** Highly Voted 4 years, 5 months ago

Should be B (3).

Incremental would require restoring Sunday's full backup, Monday's changes, and then Tuesday's changes to restore data deleted Wednesday morning.

upvoted 29 times

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

Yes. Answer is B. Sunday Full + Monday Tues Incr

upvoted 10 times

🗨️ **David_Le** 3 years, 2 months ago

You also need Tuesday's incremental backup copy.

upvoted 1 times

🗨️ **FenrirIII** Most Recent 3 years, 2 months ago

Selected Answer: B

Sun (Full) + Mon (Incr) + Tues (Incr) = 3

upvoted 3 times

🗨️ **Gruchmen** 4 years, 5 months ago

It will be 2 when there will be Differential backup stated.

Here should be 3 (B).

upvoted 5 times

Which file system uses a MapReduce framework to run applications and a NameNode for managing the file system?

- A. New Technology File System (NTFS)
- B. Hadoop Distributed File System (HDFS)
- C. Network File System (NFS)
- D. Common Internet File System (CIFS)

Suggested Answer: *B*

Reference:

<https://searchdatamanagement.techtarget.com/definition/Hadoop-Distributed-File-System-HDFS>

  **rustam_sh** 3 years, 5 months ago

correct

upvoted 3 times

When using a bridged iSCSI connection, which components facilitate communication between gateways?

- A. Compute system with iSCSI ports and the storage system with Ethernet ports.
- B. Compute system with iSCSI ports and the storage system with only FC ports.
- C. Compute system with FC HBA ports and the storage system with only FC ports.
- D. Compute system with FCoE ports and the storage system with iSCSI ports.

Suggested Answer: D

Reference:

<https://www.mycloudwiki.com/san/iscsi-interfaces/>

Community vote distribution

B (100%)

🗳️ **ameya84** Highly Voted 4 years, 5 months ago

B is correct

upvoted 13 times

🗳️ **atidd** 4 years, 5 months ago

Correct. iSCSI Gateway sits between FC SAN and IP Compute (iSCSI).

upvoted 4 times

🗳️ **mo7amed_3id** Most Recent 7 months, 1 week ago

Selected Answer: B

B is the correct answer

upvoted 1 times

🗳️ **andikovaci** 3 years, 1 month ago

Selected Answer: B

b is correct

upvoted 1 times

🗳️ **Shadyamy** 3 years, 11 months ago

Please change answer to B

upvoted 2 times

🗳️ **JBONINO** 4 years ago

Bridged iSCSI: This type of connectivity enables the initiators to exist in an IP environment while the storage systems remain in an FC SAN environment. It enables the coexistence of FC with IP by providing iSCSI-to-FC bridging functionality. The figure on the right illustrates a bridged iSCSI implementation. It shows connectivity between a compute system with an iSCSI initiator and a storage system with an FC port

upvoted 2 times

🗳️ **learner** 4 years, 3 months ago

Bridged iSCSI: This type of connectivity enables the initiators to exist in an IP environment while the storage systems remain in an FC SAN environment. It enables the coexistence of FC with IP by providing iSCSI-to-FC bridging functionality. The figure on the right illustrates a bridged iSCSI implementation. It shows connectivity between a compute system with an iSCSI initiator and a storage system with an FC port. As the storage system does not have any iSCSI port, a gateway or a multiprotocol router is used. The gateway facilitates the communication between the compute system with iSCSI ports and the storage system with only FC ports. The gateway converts IP packets to FC frames and conversely, thus bridging the connectivity between the IP and FC environments. The gateway contains both FC and Ethernet ports to facilitate the communication between the FC and the IP environments. The iSCSI initiator is configured with the gateway's IP address as its target destination. On the other side, the gateway is configured as an FC initiator to the storage system.

B is correct

upvoted 4 times

What accurately describes hypervisor-based continuous data protection?

- A. Virtual appliance is running on a hypervisor.
- B. Journal is stored in a virtual machine.
- C. Virtual appliance runs on a storage system.
- D. Replica volumes are created in the hypervisor.

Suggested Answer: A

 **Kate321** Highly Voted 3 years, 11 months ago

write answer is A.

. In this

deployment, the specialized hardwarebased appliance is replaced with virtual appliance which is running on a hypervisor. T

upvoted 5 times

 **visgoy** Most Recent 3 years, 12 months ago

Hypervisor-based CDP

The illustration shows a CDP local replication implementation.

▣ Protects a single or multiple VMs locally or remotely

▣ Enables to restore VM to any PIT

▣ Virtual appliance is running on a hypervisor

▣ Write splitter is embedded in the hypervisor

Notes


Some vendors offer continuous data protection for VMs through hypervisor-based CDP implementation. In this deployment, the specialized hardware-based appliance is replaced with virtual appliance which is running on a hypervisor. The write splitter is embedded in the hypervisor. This option protects single or multiple VMs locally or remotely and enables to restore VMs to any PIT. The local and remote replication operations are as similar as network-based CDP replication.

upvoted 1 times

 **JBONINO** 4 years ago

NO, correcta A

upvoted 3 times

 **eddb** 4 years, 1 month ago

B.

Journal is stored in a virtual machine.

upvoted 3 times

When implementing a full mesh topology in a Fibre Channel (FC) SAN environment, what is the effect when there is an increase in the number of switches?

- A. Reduces the available node ports for ISL connectivity.
- B. Increases the load on the ISLs due to traffic aggregation.
- C. Increases the hop count for compute to storage traffic.
- D. Reduces the available switch ports for node connectivity.

Suggested Answer: D

🗨️ **runnerbean** 3 years, 3 months ago

Answer D - Reduces the available switch ports for node connectivity

Page 329 ISMv4

upvoted 2 times

🗨️ **Ivan_Miran** 3 years, 3 months ago

Answer D

"5.10.1 Mesh Topology

A mesh topology may be one of the two types: full mesh or partial mesh.

In a full mesh, every switch is connected to every other switch in the topology. A full mesh topology may be appropriate when the number of switches involved is small. A typical deployment would involve up to four switches or directors, with each of them servicing highly localized host-to-storage traffic. In a full mesh topology, a maximum of one ISL or hop is required for host-to-storage traffic. However, with the increase in the number of switches, the number of switch ports used for ISL also increases. <<This reduces the available switch ports for node connectivity.>>"

Information Storage and Management . Wiley. Kindle Edition.

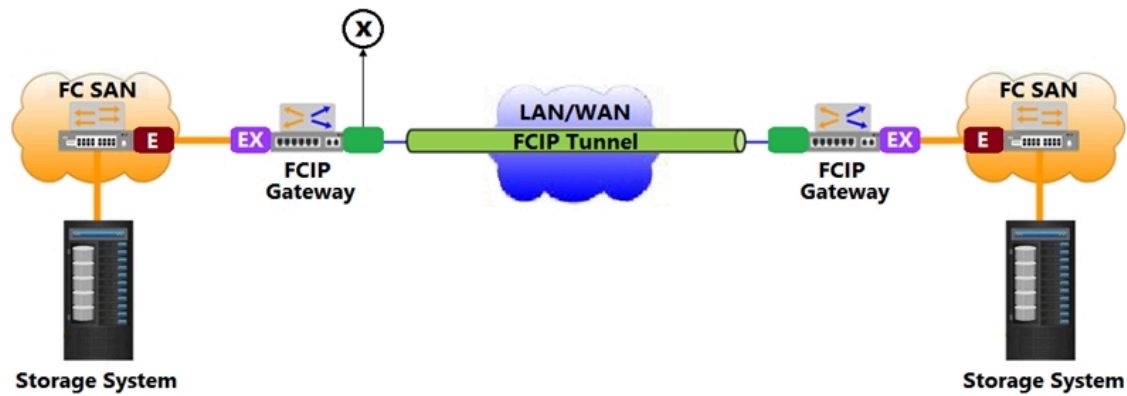
upvoted 1 times

🗨️ **JBONINO** 4 years ago

In a full mesh, every switch is connected to every other switch in the topology. A full mesh topology may be appropriate when the number of switches that are involved is small. A typical deployment would involve up to four switches or directors, with each of them servicing highly localized compute-to-storage traffic. In a full mesh topology, a maximum of one ISL or hop is required for compute-to-storage traffic. However, with the increase in the number of switches, the number of switch ports that are used for ISL also increases. This process reduces the available switch ports for node connectivity. - D correct

upvoted 4 times

Based on the exhibit, what does the "X" represent in the FCIP tunnel configuration?



- A. VX_port
- B. EX_port
- C. F_port
- D. VE_port

Suggested Answer: *D*

🗨️ **Expert2000** 3 years, 3 months ago

The FCIP gateway is also configured with a VE_port that behaves like an E_Port, except that the VE_Port is used to transport data through an FCIP tunnel

upvoted 2 times

🗨️ **grasim71** 3 years, 4 months ago

VE_port for merged and separate fabric configuration(page 390-392 in guide).

The FCIP gateway is also configured with a VE_port that behaves like an E_Port, except that the VE_Port is used to transport data through an FCIP tunnel

▮ The FCIP tunnel has VE_Ports on both ends

▮ The VE_Ports establish virtual ISLs through the FCIP tunnel, which enable fabrics on either side of the tunnel to merge

upvoted 3 times

🗨️ **no_reason** 3 years, 5 months ago

Cancel last comment..

upvoted 1 times

🗨️ **no_reason** 3 years, 5 months ago

Correct answer is A (E_Port)

upvoted 1 times

What is an accurate statement about stretched VLANs?

- A. Reduces network bandwidth requirements during VM migration across sites.
- B. Requires a change in the VMs' IP addresses to match the addressing scheme at the other sites.
- C. Enables automatic modification of VLAN addresses during VM migration across sites.
- D. Supports migration of VMs across sites without the need to change their network configurations.

Suggested Answer: A


Community vote distribution

D (100%)

  **mathmad** Highly Voted 4 years, 7 months ago

it should be D

upvoted 21 times

  **Jamestyler22** Highly Voted 4 years, 4 months ago

Stretched VLANs also enable the movement of virtual machines (VMs) between sites without the need to change their network configurations.
correct is D


upvoted 12 times

  **David_Le** Most Recent 3 years, 2 months ago

Selected Answer: D

Stretched VLANs also enable the movement of virtual machines (VMs) between sites without the need to change their network configurations.
This simplifies the creation of high-availability clusters, VM migration, and application and workload mobility across sites.



upvoted 1 times

  **rachid1761960** 3 years, 6 months ago

The correct is D

Stretched VLANs also enable the movement of virtual machines (VMs) between sites without the need to change their network configurations.
This simplifies the creation of high-availability clusters, VM migration, and application and workload mobility across sites.

upvoted 2 times

  **suzcz** 3 years, 10 months ago

D is the ans. Enables movement of VM between sites without changing their network configs

upvoted 3 times

  **fer162** 4 years, 4 months ago

from the guide "Stretched VLANs also enable the movement of virtual machines (VMs) between sites without the need to change their network configurations."

upvoted 7 times