



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



CERTIFICATION TEST

- CertificationTest.net - Cheap & Quality Resources With Best Support

Which of the following is typical of software licensing in the cloud?

- A. Per socket
- B. Perpetual
- C. Subscription-based
- D. Site-based

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **nixonbii** Highly Voted 👍 2 years, 4 months ago

Typically pay-as-you-go for a learning period, then subscription based for production.
upvoted 5 times

🗳️ 👤 **mamedov_elvin** Most Recent 🕒 1 month, 3 weeks ago

Selected Answer: C

Correct answer is C
upvoted 1 times

🗳️ 👤 **fluke92** 7 months, 1 week ago

Selected Answer: C

In cloud environments, software licensing is typically subscription-based, where users pay a recurring fee (monthly, annually, etc.) for access to the software and services. This model aligns with the on-demand nature of cloud computing
upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 2 years, 5 months ago

Selected Answer: C

Concur with C
upvoted 1 times

🗳️ 👤 **Cyduk45** 2 years, 10 months ago

Selected Answer: C

answer is C
upvoted 2 times

🗳️ 👤 **ziadatacloud** 3 years, 1 month ago

Selected Answer: C

C is the correct answer
upvoted 1 times

🗳️ 👤 **Crimson** 3 years, 9 months ago

Answer is C
upvoted 4 times

A server administrator wants to run a performance monitor for optimal system utilization. Which of the following metrics can the administrator use for monitoring?

(Choose two.)

- A. Memory
- B. Page file
- C. Services
- D. Application
- E. CPU
- F. Heartbeat

Suggested Answer: AE

Community vote distribution

AE (100%)

🗲️ 👤 **Sweety_Certified7** 11 months, 2 weeks ago

Selected Answer: AE

Obviously correct!

upvoted 1 times

🗲️ 👤 **Obi_Wan_Jacoby** 1 year, 11 months ago

Selected Answer: AE

Concur with A&E

upvoted 1 times

🗲️ 👤 **szl0144** 2 years, 3 months ago

Selected Answer: AE

AE correct without doubts

upvoted 1 times

🗲️ 👤 **Dion79** 2 years, 10 months ago

A&E - I'd go with the provided answers.

Baselines, and the related performance monitoring, begin with the four major subsystems of the server: processor, memory, storage, and network.

Processor% Processor time

Processor—% User time

Memory—Pages/sec

Network Interface—Transfers/sec

Physical Disk—Disk Transfers/sec

Physical Disk—Average disk queue length

Reference: Introduction: CompTIA CertMaster Learn - Manage Documentation Topic 4B -Manage Documentation

upvoted 2 times

After configuring IP networking on a newly commissioned server, a server administrator installs a straight-through network cable from the patch panel to the switch. The administrator then returns to the server to test network connectivity using the ping command. The partial output of the ping and ipconfig commands are displayed below:

```
ipconfig/all

IPv4 address: 192.168.1.5
Subnet mask: 255.255.255.0
Default gateway: 192.168.1.2

Pinging 192.168.1.2 with 32 bytes of data:

Reply from 192.168.1.2: Request timed out
Reply from 192.168.1.2: Request timed out
Reply from 192.168.1.2: Request timed out
Reply from 192.168.1.2: Request timed out
```

The administrator returns to the switch and notices an amber link light on the port where the server is connected. Which of the following is the MOST likely reason for the lack of network connectivity?

- A. Network port security
- B. An improper VLAN configuration
- C. A misconfigured DHCP server
- D. A misconfigured NIC on the server

Suggested Answer: D

Community vote distribution

D (100%)

🗨️ **lightthetorch** 2 weeks, 6 days ago

Selected Answer: B

Decided to shoot this one over to chatgpt, this is what it says in summary.

Both, B. Improper VLAN configuration and D. Misconfigured NIC are plausible answers depending on the assumptions:

If we assume modern hardware, VLAN misconfiguration is more likely.

If we assume older or manually configured hardware, NIC misconfiguration (e.g., cable type, duplex mismatch) becomes more likely.

If the question is from CompTIA Server+, which often assumes modern enterprise environments, B. Improper VLAN configuration is still the best answer—but D is a strong contender and worth remembering for real-world troubleshooting.

upvoted 1 times

🗨️ **PaytoPlay** 2 months ago

Selected Answer: B

B is the answer here, layer 1 and 2 are good but its not routing. This in conjunction with the amber light indicates VLAN issue.

upvoted 1 times

🗨️ **fluke92** 7 months, 1 week ago

Selected Answer: D

Given the amber link light observed on the switch port, this strongly supports the conclusion that B. An improper VLAN configuration is the most likely cause.

The VLAN mismatch would prevent the server's traffic from being routed correctly within the network, leading to the ping failures.

upvoted 1 times

🗨️ **Ronn_Burgandy** 9 months ago

Selected Answer: D

The amber light on the switch could mean A or D and they are both totally viable, both would also return request timed out. Of course this is a typical CompTIA what is the most correct answer. I think they are eluding to that the NIC may not have auto-mdi/mdix enabled because the questions mentions a straight through cable rather than a crossover cable. Even though most modern servers and switches have this enabled by default. There also isn't anything to suggest that port security would be tripped like another device having been plugged into port prior.

upvoted 1 times

🗨️ **Sweety_Certified7** 1 year, 4 months ago

Selected Answer: D

Answer is D for sure.

Check the third question on this website that has correct and verified answers (I have been using it since long): <https://quizlet.com/866168501/sk0-005-server-dump-part-1-flash-cards/>

upvoted 2 times

🗨️ **Obi_Wan_Jacoby** 2 years, 5 months ago

Selected Answer: D

I believe with options A, B and C you would get "Destination host unreachable" due to the fact there would be no way for the client/server to reach whatever they are pinging. But with D you could get request timed out due to a number of things.

<https://www.quora.com/What-is-the-difference-between-request-timed-out-and-destination-host-unreachable-in-terms-of-device-presence-using-ping>

upvoted 2 times

🗨️ **Pongsathorn** 2 years, 7 months ago

Selected Answer: D

Some devices, such as network switches, use colors to indicate performance information rather than status information. For example, a network switch may display a green LED on a port if the connection is operating at 1 Gbps but an amber LED if the port is operating at 100 Mbps. Check the device's documentation to interpret the colors correctly.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 102.

upvoted 1 times

🗨️ **dcdc1000** 2 years, 9 months ago

Hi, I'm going with A. WRT the questions, there is no reason to believe answer D i.e. the NIC is misconfigured due, the question indicates it has an IP address, DG, and appropriate NIC. Now then, if you connect the server's ethernet cable to the port and the switchport has port-security enabled, you will absolutely get a solid amber light and the result is Request timed out.

upvoted 2 times

🗨️ **PEsty93** 3 years, 2 months ago

Selected Answer: D

Amber light indicates slow speed. Straight through cable is different and only has half the pairs. The issue would be because the NIC is not running the correct Duplex. The Duplex needs to be the same at both ends.

upvoted 3 times

🗨️ **dnc1981** 3 years, 3 months ago

If the led is solid amber, then the answer is port security.

A is the correct answer.

upvoted 3 times

🗨️ **Dion79** 3 years, 2 months ago

You don't think it could be D? Why? I'm not saying it can't be Network port security, could be STP issue. Why not D first then A?

<https://www.omniseu.com/cisco-certified-network-associate-ccna/spanning-tree-port-states.php>

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960x/hardware/installation/guide/b_c2960x_hig/b_c2960x_hig_chapter_01.html

<https://community.cisco.com/t5/switching/switch-port-colour-codes/td-p/1239919>

upvoted 1 times

🗨️ **dnc1981** 3 years, 2 months ago

If it was STP that would not be a misconfigured NIC. STP is configured on the switch, not on the NIC. Also based on the info in the question, there is no reference to the NIC being configured. And the IP address, Subnet mask, and default gateway all look like they're all in the same subnet. So the IP config isn't misconfigured

upvoted 2 times

🗨️ **Dion79** 3 years, 2 months ago



Reason why I will stick with D. Yes... Network Port security STP is a feature that's what I was saying, not on the NIC. 1. Normally port would not show any lights if port was disabled using network port security. 2. ping reply would be host unreachable.

Misconfigured NIC - 1. Ping can reply with request timed out. 2. link light can be amber/Green. 3. New Server setups - mistakes are made and re-evaluate static NIC settings (Default Gateway address is incorrect?).

A response from a ping command that results in a message stating "Request timed out" occurs when there is an active network link but with no destination host to reach that can reply, since the host may be in another VLAN.

Reference: CompTIA CertMaster Practice Server+ SK0-005

upvoted 1 times

  **i_bird** 2 years, 10 months ago

@Dion79, Even your third info link suggests the answer is A.

When a Port LED is Amber the Port is not forwarding. This is possibly due to

- 1: The port being disabled by management,
- 2: an address violation, or Spanning Tree Protocol (STP).

After the port is reconfigured, the port LED can remain amber for up to 30 seconds while STP checks the switch for possible loops.

<https://networkengineering.stackexchange.com/questions/14905/why-only-one-amber-light-between-two-switches>

upvoted 2 times

A user cannot save large files to a directory on a Linux server that was accepting smaller files a few minutes ago. Which of the following commands should a technician use to identify the issue?

- A. pvdisplay
- B. mount
- C. df -h
- D. fdisk -l

Suggested Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **Pongsathorn** 7 months, 2 weeks ago

Selected Answer: C

df -h Displays storage capacity for filesystems or partitions

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 123.

upvoted 2 times

🗨️ 👤 **Drewid91** 1 year, 2 months ago

Selected Answer: C

C appears to be correct. Command will display file sizes in human readable format.

Ref: <https://www.linuxteck.com/df-command-in-linux-with-examples/>

upvoted 1 times

Following a recent power outage, a server in the datacenter has been constantly going offline and losing its configuration. Users have been experiencing access issues while using the application on the server. The server technician notices the date and time are incorrect when the server is online. All other servers are working. Which of the following would MOST likely cause this issue? (Choose two.)

- A. The server has a faulty power supply
- B. The server has a CMOS battery failure
- C. The server requires OS updates
- D. The server has a malfunctioning LED panel
- E. The servers do not have NTP configured
- F. The time synchronization service is disabled on the servers

Suggested Answer: AB

Community vote distribution

AB (100%)

🗳️ **ompk** 9 months, 2 weeks ago

Guys i think it's B and F. explanation below

B. faulty CMOS battery: The CMOS battery is a small battery on the motherboard that powers the BIOS settings and keeps track of the date and time when the server is powered off. If the CMOS battery fails, the server will lose its configuration and display an incorrect date and time when it is powered on. This can cause access issues for users and applications that rely on accurate time stamps.

F. The time synchronization service is a service that synchronizes the system clock with a reliable external time source, such as a network time protocol (NTP) server. If the time synchronization service is disabled on the servers, they will not be able to update their clocks automatically and may drift out of sync with each other and with the network. This can also cause access issues for users and applications that require consistent and accurate time across the network.

upvoted 1 times

🗳️ **Sweety_Certified7** 1 year, 4 months ago

Selected Answer: AB

Answers A and B are surely the answers.

Check this website with correct and verified answers: <https://quizlet.com/866168501/sk0-005-server-dump-part-1-flash-cards/> (check the 5th question)

upvoted 1 times

🗳️ **Grumpy_Old_Coot** 1 year, 7 months ago

A network remediation VLAN would explain the going offline (not -rebooting-) bit, as the server would get kicked to another IP#. And NTP issues in general will cause SSL and TLS to fail.

upvoted 1 times

🗳️ **Dingos** 2 years, 3 months ago

Chat GPT thinks it is B & E.

Asked him about E that it is trick question (only plural in answers).

Here it its answer on that topic (I do not agree, I would go with A and B):

You are correct that there is a discrepancy between the singular form of the server mentioned in the question and the plural form used in answer option E.

It's possible that the discrepancy is intentional, but it could also be a mistake. However, even if the answer option E refers to servers in general, it is still a valid option since lack of NTP configuration on any of the servers could lead to time synchronization issues and access problems for users.

In any case, it's always a good practice to carefully read and analyze the question and all answer options before selecting an answer.

upvoted 1 times

🗳️ **anywayz70** 2 years, 3 months ago

Has anyone taken the exam lately to confirm what the actual answer is?

upvoted 2 times

🗳️ 👤 **papahawaii** 2 years, 5 months ago

Look at F very closely. "The time synchronization service is disabled on the servers" <-- plural. The issue is only occurring on a single server in the question, so why would "servers" be an answer.

I chose A due to going offline, and B as the time is incorrect. Originally thought well E but it also states "servers -- plural". So only A and B really fit in this question.

upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 2 years, 5 months ago

Selected Answer: AB

I concur with A&B

upvoted 1 times

🗳️ 👤 **Pongsathorn** 2 years, 7 months ago

Selected Answer: AB

datacenter has been constantly going offline = the power supply no longer functions properly since the power returned.

losing its configuration = CMOS has no battery to reserve configuration.

upvoted 3 times

🗳️ 👤 **nixonbii** 2 years, 8 months ago

A and B are the very first places to check. This would be especially important because the question mentions a power outage.

upvoted 2 times

🗳️ 👤 **paperburn** 2 years, 8 months ago

A B are correct

upvoted 2 times

🗳️ 👤 **szl0144** 2 years, 9 months ago

Selected Answer: AB

A B are correct

upvoted 2 times

🗳️ 👤 **Dion79** 3 years ago

Selected Answer: AB

Asking for server not servers.. CMOS and PSU would cause the current issue and also wording is misleading. I would go with A and B

upvoted 3 times

🗳️ 👤 **dnc1981** 3 years, 3 months ago

A and B are the correct answers

upvoted 3 times

🗳️ 👤 **Dion79** 3 years, 3 months ago

This answer looks incorrect to me. I would select B & F

The complementary metal-oxide-semiconductor (CMOS) chip holds basic configuration information for the server. It has a replaceable battery that allows the settings to save while power is disconnected. The battery will eventually fail. When that occurs, settings stored in the basic input/output system (BIOS) on the CMOS may be reset. The date and time may also not display correctly in the BIOS.

upvoted 3 times

🗳️ 👤 **Dion79** 3 years, 2 months ago

I'm way off again by a word and sentence structure, probably not F. Lean towards A.. Like DNC1981


upvoted 1 times

A company has implemented a requirement to encrypt all the hard drives on its servers as part of a data loss prevention strategy. Which of the following should the company also perform as a data loss prevention method?

- A. Encrypt all network traffic
- B. Implement MFA on all the servers with encrypted data
- C. Block the servers from using an encrypted USB
- D. Implement port security on the switches

Suggested Answer: A

Community vote distribution



🗳️ 👤 **Kabbath1986** 1 month, 1 week ago

Selected Answer: B

The correct answer is: B. Implement MFA on all the servers with encrypted data

📄 Why this is correct:

Encrypting hard drives protects data at rest, but if an attacker or unauthorized user gains access to the server (e.g., via stolen credentials), they could still access decrypted data.

Multi-Factor Authentication (MFA) adds a strong layer of access control – even if credentials are compromised, the attacker still needs a second factor (like a phone app or hardware token).

This aligns well with a data loss prevention (DLP) strategy by preventing unauthorized access to encrypted data.

upvoted 1 times

🗳️ 👤 **fluke92** 7 months, 1 week ago

Selected Answer: B

Chat GPT:

While encrypting hard drives is an excellent data loss prevention (DLP) strategy, it is equally important to secure access to the servers where encrypted data is stored. Implementing Multi-Factor Authentication (MFA) ensures that even if an unauthorized user gains access to credentials, they will still need a second authentication factor (e.g., a mobile app code, hardware token, or biometric data) to access the server. This adds an essential layer of security to protect sensitive data.

upvoted 2 times

🗳️ 👤 **SecNoob27639** 9 months, 1 week ago

Selected Answer: A

A is the only real DLP option presented. B is access control which is part of DLP, but the big-3 of DLP are protecting Data at Rest (server), Data in Transit (network) and Data in Use (user-end). C and D are also more about access control and network segmentation than they are about preventing data loss.

upvoted 1 times

🗳️ 👤 **c32afa8** 10 months ago

Port security is a vital component of network security that helps protect network ports from unauthorized access and potential security threats. By implementing robust port security measures, organizations can safeguard sensitive data, maintain network integrity, and ensure compliance with regulatory requirements. Port security is the more correct answer where it is giving you an answer. A is correct, but vague. The question is hinting at Data in transit, as it is giving you the other end of the scenario already.

upvoted 1 times

🗳️ 👤 **Fart2023** 11 months, 3 weeks ago

Selected Answer: B

Another BS CompTIA question....

upvoted 3 times

🗳️ 👤 **Sweety_Certified7** 1 year, 4 months ago

Selected Answer: A

Answer A

upvoted 1 times

🗳️ 👤 **GRIN13** 1 year, 6 months ago

Selected Answer: B

B is correct because MFA is the next step to prevent data loss Prevention.

upvoted 2 times

🗳️ 👤 **Grumpy_Old_Coot** 1 year, 9 months ago

Selected Answer: C

We're looking at Data Loss Prevention here - which is "removing" data from where it is supposed to stay, not accessing data when we are not supposed to - Blocking USB media (Flash drives, thumb drives, external hard drives, etc) would prevent copying ("removing") data from the network.

upvoted 2 times

🗳️ 👤 **kloug** 2 years, 3 months ago

no aaaa corrcet

upvoted 1 times

🗳️ 👤 **kloug** 2 years, 3 months ago

bbbbbbbbbbbbbbbb

upvoted 2 times

🗳️ 👤 **Obi_Wan_Jacoby** 2 years, 5 months ago

Selected Answer: D

I believe D is it. Of the options, this is the only one that strikes me as a method. Study up the differences in strategy vs method

upvoted 1 times

🗳️ 👤 **Pongsathorn** 2 years, 7 months ago

Selected Answer: A

There are two different times when data encryption protects information: when the data is in transit across the network or at rest on the drive.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 214.

You have implemented "Data at rest" and then you should do "Data in transit".

upvoted 3 times

🗳️ 👤 **TylerKiro** 2 years, 8 months ago

Selected Answer: D

I wouldn't choose 'A' because not all network traffic needs to be encrypted, Client-Client or Traffic over the internet encryption seems wild. No need for 'B' since it is already encrypted. 'C'... If you do that then no USBs are usable and that's needed for a lot of things. Which leaves 'D' to be the most logical answer.

upvoted 1 times

🗳️ 👤 **paperburn** 2 years, 8 months ago

You should not have physical access to the ports so I would look at data in motion (network traffic)

upvoted 1 times

🗳️ 👤 **King2** 2 years, 9 months ago

Selected Answer: A

Reviewed answers again and took back my last answer (C)

Answer seems to be A.

The company encrypted data "at rest", so they need to encrypt data "in transit".

upvoted 2 times

🗳️ 👤 **King2** 2 years, 9 months ago

Selected Answer: C

I think C is correct.

upvoted 2 times



🗳️ 👤 **dnc1981** 3 years, 2 months ago

I think the correct answer might be A or C. If data is to be encrypted at rest, it probably should be encrypted at rest as well. Or you might block USB drives to avoid data being exfiltrated. Even if the USB drives are encrypted they still represent a risk because data could be exfiltrated via USB drive

upvoted 3 times

🗳️ 👤 **Dion79** 3 years, 2 months ago

I agree with you. I'd probably go with C.
upvoted 1 times

  **szl0144** 2 years, 9 months ago
I will go with C
upvoted 1 times

A systems administrator is setting up a server on a LAN that uses an address space that follows the RFC 1918 standard. Which of the following IP addresses should the administrator use to be in compliance with the standard?

- A. 11.251.196.241
- B. 171.245.198.241
- C. 172.16.19.241
- D. 193.168.145.241

Suggested Answer: C

Reference:

<https://whatis.techtarget.com/definition/RFC-1918>

Community vote distribution

C (100%)

🗨️ 👤 **gingasaurusrex** 11 months, 2 weeks ago

RFC 1918 was used to create the standards by which networking equipment assigns IP addresses in a private network. A private network can use a single public IP address. The RFC reserves the following ranges of IP addresses that cannot be routed on the Internet:

10.0.0.0 - 10.255.255.255 (10/8 prefix)

172.16.0.0 - 172.31.255.255 (172.16/12 prefix)

192.168.0.0 - 192.168.255.255 (192.168/16 prefix)

upvoted 3 times

🗨️ 👤 **Obi_Wan_Jacoby** 11 months, 3 weeks ago

Selected Answer: C

I concur with C

upvoted 1 times

🗨️ 👤 **Ariel235788** 1 year, 10 months ago

One thing that stood out to me from a context perspective, C is the only private IP address. Could only assume this was referring to private IP addressing haha


upvoted 3 times

An administrator needs to perform bare-metal maintenance on a server in a remote datacenter. Which of the following should the administrator use to access the server's console?

- A. IP KVM
- B. VNC
- C. A crash cart
- D. RDP
- E. SSH

Suggested Answer: A

Community vote distribution



49tiktok Highly Voted 2 years, 9 months ago

Selected Answer: A

I disagree. If a system is bare metal, there is no way to access the BIOS and Hardware without a crash cart, KVM, console, or remote management (iLO, etc). Since only one of those options (ip KVM or iLO) can be done remotely, that would be the correct answer.

RDP would allow you to manage the operating system running on the server, but not the bare metal system itself.

upvoted 6 times

Kraken84 1 year, 8 months ago

But it says he is in the remote data center. So if he is in fact inside the data center, wouldn't a crash cart make sense?

upvoted 2 times

c32afa8 Most Recent 10 months ago

First off, RDP is only a Windows based protocol, not all servers use Windows as the OS.

At one time, all servers were bare-metal servers. Servers were kept on-premises and often belonged to the organization using and operating them. In computer networking, a bare-metal server is a physical computer server that is used by one consumer, or tenant, only. Each server offered for rental is a distinct physical piece of hardware that is a functional server on its own. They are not virtual servers running in multiple pieces of shared hardware. With that being the definition of On-Premises Server, you could use A. IP KVM, or C. A crash cart to gain access to the server. The question is vague enough not to be able to correctly determine by not knowing the need. Is the server in the same room? Do they have a OOB network set up? If you can access the server via IP KVM, why wouldn't you opt to?. Make sure you do your Backups/ Testing beforehand.

upvoted 1 times

Kraken84 1 year, 8 months ago

Selected Answer: C

An administrator needs to perform bare-metal maintenance on a server "IN" a remote datacenter.

upvoted 1 times

Obi_Wan_Jacoby 2 years, 5 months ago

Selected Answer: A

I concur with A. 49tiktok has the correct reasoning as to why

upvoted 1 times

lordguck 2 years, 7 months ago

A: This way, the OS on the server does not matter and you have access to BIOS and other functions like RAID configuration.

upvoted 2 times

Pongsathorn 2 years, 8 months ago

Selected Answer: A

KVM vendors have responded to the need for a KVM switch that can be accessed over the network. The switch is like the one you saw earlier with one difference—it can be reached through the network. This means it is accessible not only from a workstation in the next room, but from anywhere. In this particular implementation (it can be done several ways), each server has a small device between it and the KVM switch that accepts the serial and keyboard/mouse connections.

upvoted 3 times

  **paperburn** 2 years, 8 months ago

KVM-over-IP switches are the remote access version of local KVM (keyboard/video/mouse) devices. They allow administrators to control a remote system by sending keyboard and mouse signals over the network and displaying the remote system's video output on the administrator's local machine via a web browser. Administrators can perform almost any function on the remote system as long as it's powered on and connected to the network.


upvoted 2 times

A technician needs to provide a VM with high availability. Which of the following actions should the technician take to complete this task as efficiently as possible?

- A. Take a snapshot of the original VM
- B. Clone the original VM
- C. Convert the original VM to use dynamic disks
- D. Perform a P2V of the original VM

Suggested Answer: B

Community vote distribution



Obi_Wan_Jacoby Highly Voted 2 years, 5 months ago

Selected Answer: B

I believe the answer to be B. The answer A (Snapshot) does not fall under the classification for high availability. For high availability you need a cluster. You would clone a server and place that clone within a cluster for fail-over (fail-over cluster)

upvoted 8 times

fluke92 Most Recent 7 months, 1 week ago

Selected Answer: B

To achieve high availability, the technician needs to ensure that the virtual machine (VM) can recover or failover quickly in the event of a failure. Cloning the original VM is an efficient way to create an identical backup that can be quickly brought online or used for failover purposes.

In high-availability setups, the cloned VM can be deployed on a separate host or in a failover cluster to ensure continuous availability.

upvoted 1 times

kloug 2 years, 3 months ago

bbbbbbbbbb

upvoted 1 times

Pongsathorn 2 years, 7 months ago

Selected Answer: A

The keyword is "efficiently", so the answer is Snapshot.

upvoted 3 times

FreePrivacy 11 months, 2 weeks ago

Snapshot, doesn't provide high availability, to meet both high availability and efficiency, the answer is B.

upvoted 1 times

A server administrator receives a report that Ann, a new user, is unable to save a file to her home directory on a server. The administrator checks Ann's home directory permissions and discovers the following: dr-xr-xr-- /home/Ann

Which of the following commands should the administrator use to resolve the issue without granting unnecessary permissions?

- A. `chmod 777 /home/Ann`
- B. `chmod 666 /home/Ann`
- C. `chmod 711 /home/Ann`
- D. `chmod 754 /home/Ann`

Suggested Answer: D

Reference:

<https://linuxize.com/post/what-does-chmod-777-mean/>

Output

```
-rw-r--r-- 12 linuxize users 12.0K Apr  8 20:51 filename.txt
|[-][-][-]-  [-----] [---]
| | | | |      |      |
| | | | |      |      +-----> 7. Group
| | | | |      +-----> 6. Owner
| | | | +-----> 5. Alternate Access Method
| | | +-----> 4. Others Permissions
| | +-----> 3. Group Permissions
| +-----> 2. Owner Permissions
+-----> 1. File Type
```

Community vote distribution

D (100%)

Obi_Wan_Jacoby 11 months, 3 weeks ago

Selected Answer: D

I concur with D

upvoted 1 times

azre_certified1111 1 year ago

The answer should be `chmod 554 /home/Ann`

upvoted 1 times

azre_certified1111 1 year ago

My mistake. I misread the question


upvoted 1 times

Which of the following documents would be useful when trying to restore IT infrastructure operations after a non-planned interruption?

- A. Service-level agreement
- B. Disaster recovery plan
- C. Business impact analysis
- D. Business continuity plan

Suggested Answer: B

Community vote distribution



slpcomputer Highly Voted 3 years, 5 months ago

I too believe it should be the DRP...BIA is the planning stage that establishes the objectives that go into the DRP.
upvoted 8 times

nixonbii Most Recent 9 months ago

Selected Answer: B

"business impact analysis (BIA) - Identifies the degree and scope of impact on a business when negative incidents occur. The BIA enables easier prioritization of assets that require protection from threats." - CompTia Server+ Certification Exam Guide, McGraw Hill 2nd ed. SK0-005. The book does not indicate that it should be consulted in the course of recovering from an incident, rather it stresses the importance of having it in place before an incident occurs. I vote for DRP.
upvoted 2 times

Pongsathorn 9 months ago

Selected Answer: B

B is correct.

A disaster recovery plan (DRP) is a documented, structured approach that describes how an organization can quickly resume work after an unplanned incident. A DRP is an essential part of a business continuity plan (BCP). It is applied to the aspects of an organization that depend on a functioning information technology (IT) infrastructure. A DRP aims to help an organization resolve data loss and recover system functionality so that it can perform in the aftermath of an incident, even if it operates at a minimal level.

upvoted 1 times

Pongsathorn 9 months ago

Selected Answer: B

BCP

The purpose of a BCP is to address how to maintain business obligations in the case of a large-scale disaster (natural or caused by humans). The focus is on the business as a whole, rather than on the areas for which IT is responsible.

Business Impact Analysis (BIA) and Disaster Recovery Plans (DRPs) are subsets of the BCP.

BIA

A BIA is a component of your company's overall BCP. The purpose of the BIA is to identify the potential consequences of an interruption to your business due to a disaster or other unplanned events.

DRP

The Disaster Recovery Plan (DRP) specifies responsibilities, tools, procedures, and supporting policies for recovering IT services in the event of a disaster. While the BCP focuses on the business as a whole, the DRP focuses on IT services.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 78-79.

upvoted 2 times

error77 1 year, 3 months ago

Selected Answer: B

Question says "trying to restore", meaning service has not been resumed yet, so the most urgent thing to do is restore the service - answer is B.
upvoted 1 times

RBL23168 1 year, 4 months ago

Quite obviously B.

upvoted 1 times

🗳️ 👤 **Sweety_Certified7** 1 year, 5 months ago

Selected Answer: C

A DRP is a strategy to restore the ordinary business operations of an organization's "IT" (Rather than the whole business) infrastructure following a disaster scenario. Business continuity planning is about maintaining operations "during" a disaster, while a cloud disaster recovery plan focuses on restoring IT infrastructure and systems "after" a severe disaster.

Reference: <https://www.atlantic.net/disaster-recovery/disaster-recovery-plan/>

upvoted 3 times

🗳️ 👤 **cloudchief25** 1 year, 7 months ago

Everyone answering B are thinking about this wrong. The answer is C

Lets use an example. The company has a datacenter. You have a major power outage that last for days. Much longer then UPS are designed to support. At that point of the power outage is when you would be looking at your DRP. DRPs should lay out different scenarios depending on the extent of the outage or disaster and what you should be doing DURING a disaster.

The question is, what document would you refer to AFTER a unplanned outage. So the outage or disaster is over. At this point you need to look at what damage has the company suffered because of this outage. Hence why the answer is C.

upvoted 2 times

🗳️ 👤 **error77** 1 year, 3 months ago

Question says "trying to restore" after the outage, meaning the service has not resumed to normal yet, so answer is B.

upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 2 years, 5 months ago

Selected Answer: B

I too select B. The reason is because it specifically indicates you need to restore "IT infrastructure operations". This would describe what a DRP is for. DRP's for each dept/services/assets (in this case, for infrastructure, things like DHCP, DNS) are created under the BIA and priority given for each.

upvoted 1 times

🗳️ 👤 **paperburn** 2 years, 8 months ago

Selected Answer: D

A business continuity plan (BCP) is a document that consists of the critical information an organization needs to continue operating during an unplanned event. The BCP states the essential functions of the business, identifies which systems and processes must be sustained, and details how to maintain them.

upvoted 1 times

🗳️ 👤 **paperburn** 2 years, 8 months ago

A business continuity plan (BCP) is a document that consists of the critical information an organization needs to continue operating during an unplanned event. The BCP states the essential functions of the business, identifies which systems and processes must be sustained, and details how to maintain them.

upvoted 1 times

🗳️ 👤 **paperburn** 2 years, 8 months ago

A BRP will help you respond effectively if an incident or crisis affects your business. It aims to shorten your recovery time and minimize losses. Your recovery plan contains information relating to planning for recovery as well as the resumption of critical business activities after a crisis has occurred.

upvoted 1 times

🗳️ 👤 **paperburn** 2 years, 8 months ago

A business continuity plan (BCP) i sorry

upvoted 1 times

🗳️ 👤 **Drewid91** 2 years, 10 months ago

Selected Answer: B

Seems like Disaster Recovery Plan makes the most sense. Other documents might be consulted or updated after.

upvoted 4 times

🗳️ 👤 **Lajoni** 3 years, 3 months ago

C is correct because BIA evaluates the risks & recovery procedures when a disaster hits the company. This document might include a section dedicated to a recovery procedure if a change causes a problem

upvoted 2 times

🗨️ 👤 **dnc1981** 3 years, 3 months ago

DRP is the answer because the question asks about restoring AFTER an unplanned outage. DRP is the document you would consult for this. BIA would be more high level and would be written BEFORE an outage, to plan for how to recover from such an outage

upvoted 3 times

🗨️ 👤 **Rodmas** 3 years, 7 months ago

your right

upvoted 3 times

🗨️ 👤 **DayWalker144** 3 years, 9 months ago

BIA doesn't sound right to me, BIA is part of the BCP and BCP identifies critical systems and services, DRP is the execution of BCP, SLA - doesn't apply. Anyone else find this answer odd or can give more input on this? I'm leaning towards DRP, any input?

upvoted 4 times

🗨️ 👤 **Dion79** 3 years, 2 months ago

DRP is probably correct. And BIA does sound off.. provided answer is not correct.. I'd go with DRP..

upvoted 1 times

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

I do agree with you. DRA (B) should be the answer here. The text below is taken from Chapter 8: Disaster Recovery Planning in the CompTIA Server+ Certification All-in-One Exam Guide, Second Edition by Daniel Lachance

"Disaster Recovery Plan

A disaster recovery (DR) plan is used to bring failed systems online as quickly and efficiently as possible. The DR plan must be updated periodically to reflect changing threats.

The DR plan contains step-by-step procedures detailing exactly how systems are to be quickly recovered. All stakeholders must know their roles for the effective recovery of failed systems.

The DR plan includes the following:

- Table of contents
- DR scope
- Contact information
- Recovery procedures
- Document revision history
- Glossary"

upvoted 2 times

A systems administrator is setting up a new server that will be used as a DHCP server. The administrator installs the OS but is then unable to log on using Active Directory credentials. The administrator logs on using the local administrator account and verifies the server has the correct IP address, subnet mask, and default gateway. The administrator then gets on another server and can ping the new server. Which of the following is causing the issue?

- A. Port 443 is not open on the firewall
- B. The server is experiencing a downstream failure
- C. The local hosts file is blank
- D. The server is not joined to the domain

Suggested Answer: D

Community vote distribution

D (100%)

🗨️ 👤 **Obi_Wan_Jacoby** 11 months, 3 weeks ago

Selected Answer: D

Concur with D

upvoted 1 times

🗨️ 👤 **brbell6238** 1 year ago

D is the correct answer

upvoted 1 times

A systems administrator is preparing to install two servers in a single rack. The administrator is concerned that having both servers in one rack will increase the chance of power issues due to the increased load. Which of the following should the administrator implement FIRST to address the issue?

- A. Separate circuits
- B. An uninterruptible power supply
- C. Increased PDU capacity
- D. Redundant power supplies

Suggested Answer: A

Community vote distribution

A (67%)

C (33%)

🗳️ 👤 **Atemius** Highly Voted 3 years, 3 months ago

Is C correct? This question is worded so poorly. I copied the passage from Chapter 2: Server Hardware in the CompTIA Server+ Certification All-in-One Exam Guide, Second Edition by Daniel Lachance

"Power distribution units (PDUs) provide power outlets to racks in server rooms and data centers. To eliminate a single point of failure, redundant PDUs should be plugged into separate circuits. To extend this point, redundant server power supplies should each plug into separate PDUs. Data centers normally have alternate sources or providers of power, such as diesel generators, in the case of a power outage.

Because many different types of items can draw power from PDUs, you should check your PDU's rating to ensure that your equipment doesn't draw more power than the PDU's load capacity can accommodate."

upvoted 8 times

🗳️ 👤 **TheITStudent** 2 years, 10 months ago

@Atemius, thanks for posting!!! based on this, I am changing my answer to A- "To eliminate a single point of failure, redundant PDUs should be plugged into separate circuits."

upvoted 1 times

🗳️ 👤 **kx7tg4xu** Most Recent 8 months, 2 weeks ago

When you're faced with a power problem with only two additional servers, "Separate the circuits and that's the solution!" How foolish to think that.

The problem statement does not say that only one power circuit exists. But by process of elimination, what remains is "A. Separate circuits."

upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 9 months ago

Selected Answer: A

Answer "A" is the only option. Answer B could still run both servers run off the same load. Answer C is basically just adding more outlets, they too can still pull from the same source. Redundant power supplies (Answer D) should be used, but they too can pull from the same source. You need separate circuits. The redundant power supply, outlets or UPS (answers BC&D) can run from the other circuit.

upvoted 2 times

🗳️ 👤 **kx7tg4xu** 10 months, 2 weeks ago

Selected Answer: A

I think A is correct on this one.

The other options are not appropriate as a first step.

upvoted 1 times

🗳️ 👤 **Grumpy_Old_Coot** 1 year, 9 months ago

Selected Answer: C

This one's wording will trip you up. "The administrator is concerned that having both servers in one rack will increase the chance of power issues due to the increased load." Servers usually have redundant power supplies. The load handling portion of properly configured rack is usually separate circuits feeding redundant UPS units each connected to a PDU. If a PDU cannot handle the amperage...

upvoted 2 times

🗳️ 👤 **gingasaurusrex** 2 years, 4 months ago

Selected Answer: A

It has to Be A

upvoted 1 times

🗨️ 👤 **RickICT** 2 years, 5 months ago

Since there is no indication of the rack capacity prior to installing the new servers, I believe this question have two answers, B & C. If I am going to install an additional PDU, I should also ensure there is an additional UPS to connect it to which will ensure redundancy. Connecting a second PDU to a single UPS won't help with increase the capacity.

upvoted 1 times

🗨️ 👤 **bash45** 2 years, 6 months ago

Power distribution units (PDUs) provide power outlets to racks in server rooms and data centers. To eliminate a single point of failure, redundant PDUs should be plugged into separate circuits. To extend this point, redundant server power supplies should each plug into separate PDUs. Data centers normally have alternate sources or providers of power, such as diesel generators, in the case of a power outage. Because many different types of items can draw power from PDUs, you should check your PDU's rating to ensure that your equipment doesn't draw more power than the PDU's load capacity can accommodate

upvoted 3 times

🗨️ 👤 **lordguck** 2 years, 7 months ago

I vote for A. It makes sense to use two PDUs to supply each server separately with power, if the servers have two power supplies, cross connect them, one to each PDU.

upvoted 1 times

🗨️ 👤 **Timock** 2 years, 9 months ago

Question states "increase the chance of power issues due to the increased load."

" Which of the following should the administrator implement FIRST to address the issue?"

- So the worry here is increased load and how to address this FIRST.

An increased PDU for power distribution would be the first thing to do to increase the power load which should address the "issues due to increased load."

upvoted 2 times

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

It is best practice to have your PDU load capacity 50% more than your planned server load capacity on a rack

So make sure you know your planned server load capacity before server installation and increase your PDU load capacity to be atleast 50% more than the load capacity of your servers on your rack

upvoted 1 times

Which of the following is a method that is used to prevent motor vehicles from getting too close to building entrances and exits?

- A. Bollards
- B. Reflective glass
- C. Security guards
- D. Security cameras

Suggested Answer: A

Reference:

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

Community vote distribution

A (100%)

🗳️ 👤 **dinosan** 7 months, 2 weeks ago

Selected Answer: A

Bollards is the only correct option!

upvoted 1 times

🗳️ 👤 **Pongsathorn** 7 months, 2 weeks ago

Selected Answer: A

The building may also be protected by bollards that enforce a buffer zone between the building and motor traffic. These bollards help to prevent ramming attacks. Bollards are common near building entrances.

upvoted 1 times

🗳️ 👤 **Fineb** 9 months, 1 week ago

Yes, I agreed with A

upvoted 2 times

🗳️ 👤 **Atemius** 1 year, 3 months ago

Correct answer is A. Chapter 6 Server and Network Security in the CompTIA Server+ Certification All-in-One Exam Guide, Second Edition by Daniel Lachance

"The first line of physical defense is perimeter security, which comes from the following:

- Fencing
- Bollard posts to protect buildings from vehicle incursion
- Lighting
- Locked gates
- Security guards
- Guard dogs
- Limited access to areas of a facility
- Motion-sensing security systems"

upvoted 2 times

A technician is installing a variety of servers in a rack. Which of the following is the BEST course of action for the technician to take while loading the rack?

- A. Alternate the direction of the airflow
- B. Install the heaviest server at the bottom of the rack
- C. Place a UPS at the top of the rack
- D. Leave 1U of space between each server

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **Noms100** Highly Voted 3 years ago

B, Always load servers from the bottom up. This prevents servers from becoming top-heavy and tipping also we do not alternate the direction of the airflow because that would have already been setup on the room with airflow sensors
upvoted 6 times

🗳️ 👤 **Dion79** 3 years ago

Not true and not from what I learned. normally heaviest equipment is Power supply units which take up 3U and two or three of those thing can be used. Servers should be installed in the middle to top of Server Rack. Something like this: KVM console, KVM Switch, Layer2 or 3 switch depending, servers, 2u space, then PSU.
upvoted 2 times

🗳️ 👤 **skid40** Most Recent 9 months, 1 week ago

D is the Correct Answer
upvoted 1 times

🗳️ 👤 **Riseofashes** 2 years, 2 months ago

Selected Answer: B

A is clearly the joke answer.

Data centers have cold aisles and hot aisles. Cold air blows into the cold aisle for all the devices to pull through, exhausting out to the hot aisle.

Servers pull air through the front and push hot air out the back. Imagine you had two servers with alternating airflows... They'd just cycle hot air between each other!

1U spaces between servers is good if you have the space, but not essential. Putting heavy things at the top of an empty server is a recipe for disaster.
upvoted 1 times

🗳️ 👤 **kadamske** 2 years, 2 months ago

A is the correct answer, airflow is the first thing to be considered. New servers come with server rail, regardless the sizes, with the server rail and server lift, you only need to slide it out with no energy required. If you place the bigger one at the bottom and block the airflow, all the server on the rack will be overheating quickly.
upvoted 1 times

🗳️ 👤 **Pongsathorn** 2 years, 7 months ago

Selected Answer: B

It is essential to ensure to secure the rack to the floor. This prevents servers that are pulled forward on rails from overbalancing the rack and causing it to tip over. Floors must also be rated to properly support the weight of the servers, the rack, and any other IT equipment.
upvoted 1 times

🗳️ 👤 **nixonbii** 2 years, 8 months ago

Selected Answer: B

This question is far too vague to interpret without making a lot of assumptions, however, as a rule of thumb, one NEVER puts the heaviest object at the top of a vertical structure. Heavy objects should always be on the bottom in order of heaviest to lightest. I guess OSHA doesn't inspect data centers.
upvoted 3 times

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

Selected Answer: B

Always load servers from the bottom up

upvoted 2 times

🗨️ 👤 **szl0144** 2 years, 9 months ago

Selected Answer: B

B is the answer.

upvoted 2 times

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

@Removed and Dion79: Keyword say the BEST course of action..

D is an optional implementation, while B is a standard implementation to prevent top-heavy rack.

upvoted 1 times

🗨️ 👤 **Dion79** 3 years, 1 month ago

Server Racks can be completely full with no space to accommodate. So leaving 1U space is not always the correct answer. It seems that its pointing to Hot & Cold aisles but the question almost seems its missing something.

upvoted 2 times

🗨️ 👤 **PEsty93** 3 years, 2 months ago

Selected Answer: B

You don't alternate airflow. Airflow always goes front to back.

UPS doesn't go at the top as it is probably the heaviest thing.

Leaving 1U gap is wasteful and not required. You should install everything from the bottom upwards and ensure the rack is not top heavy.

therefore correct answer is to install heaviest server at the bottom.

upvoted 1 times

🗨️ 👤 **dnc1981** 3 years, 3 months ago

B or D is the right answer

upvoted 1 times

A technician is configuring a server that requires secure remote access. Which of the following ports should the technician use?

- A. 21
- B. 22
- C. 23
- D. 443

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ **Grumpy_Old_Coot** 9 months ago

Selected Answer: B

SSH. 443 is web traffic (shopping). There's server administration stuff that works over 443, but that's mostly vmWare (vSphere) and the like.
upvoted 1 times

🗳️ **Ckamunga** 11 months, 1 week ago

B. SSH is used to provide a secure remote access to a server
upvoted 2 times

🗳️ **harukaze1337** 1 year ago

Selected Answer: B

I say B
Port 22 is for secure remote administration access (SSH).
upvoted 2 times

🗳️ **junior76w62** 1 year, 3 months ago

Port 22 is used for remote ADMINISTRATION access, Port 443 is secure access to the internet, so I believe is 443
upvoted 1 times

🗳️ **Obi_Wan_Jacoby** 1 year, 5 months ago

Selected Answer: B

Concur with B
upvoted 2 times

🗳️ **RickICT** 1 year, 5 months ago

I believe this should be port 22 for SSH. It will work on both Windows and Linux servers.
upvoted 1 times

🗳️ **Pongsathorn** 1 year, 7 months ago

Selected Answer: B

Maintaining server functionality requires the ability to connect to the server and gather information. As you learned earlier, there are many ways of connecting to the server to issue commands or display information.
SSH—this is a standard tool to connect to Linux servers and network devices. It is not commonly used with Windows, but it can be added. SSH is a secure method for remotely administering your servers.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 199.
upvoted 1 times

🗳️ **Pongsathorn** 1 year, 7 months ago

Selected Answer: B



The default port for SSH client connections is 22.
By default, these two protocols are on their standard port number of 80 for HTTP and 443 for HTTPS.
upvoted 2 times

🗳️ **nixonbii** 1 year, 8 months ago

Selected Answer: B



Have to stick with B. ports 22 and 3389 (Microsoft RDP) are staples of remote access and maintenance. 443 is best left to secure web traffic.

upvoted 2 times

  **paperburn** 1 year, 8 months ago

443 also works for RDP

upvoted 1 times

  **jagoichi** 1 year, 9 months ago

SSH port 22 for secure remote access443 is https for encrypted website browsing

upvoted 3 times

A server administrator is using remote access to update a server. The administrator notices numerous error messages when using YUM to update the applications on a server. Which of the following should the administrator check FIRST?

- A. Network connectivity on the server
- B. LVM status on the server
- C. Disk space in the /var directory
- D. YUM dependencies

Suggested Answer: C

Community vote distribution

C (71%)

D (29%)

🗳️ 👤 **dnc1981** Highly Voted 3 years, 3 months ago

Check network connectivity first
upvoted 5 times

🗳️ 👤 **Dion79** 3 years ago

Wouldn't the Tech already have network connectivity via remote connection? So if he's already connected to the server remotely then he has network connectivity. Now I would check Dependencies issues with YUM. Just thinking outloud.
upvoted 10 times

🗳️ 👤 **Rainkoot** 2 years, 10 months ago

If YUM is operations but just getting errors wouldn't the dependencies already be satisfied? The link I posted is from 9 years ago and is a situation where C is correct
upvoted 2 times

🗳️ 👤 **badgerino** Most Recent 9 months ago

Selected Answer: C

Not enough info, what are the error messages? Classic vague CompTIA question.

YUM is designed to automatically install linux package dependencies so its not D.

The admin is remotely connected, so it has network connectivity, does not state 'internet' connectivity so its not A.

B makes no sense at all.

Most likely C just from process of elimination I'm guessing, but its CompTIA so who knows what they think is 'correct'.

upvoted 1 times

🗳️ 👤 **kx7tg4xu** 10 months, 2 weeks ago

Selected Answer: C

C is the correct answer to this question.

It is efficient to start with the more basic and most likely root cause of the problem as the first point to check. Checking disk capacity is a good first step because it is easy and quick.

upvoted 1 times

🗳️ 👤 **kloug** 2 years, 3 months ago

aaaaaaaaaa

upvoted 1 times

🗳️ 👤 **szl0144** 2 years, 9 months ago

Selected Answer: D

I believe D is correct.

upvoted 2 times

🗳️ 👤 **Rainkoot** 2 years, 10 months ago

Selected Answer: C

<https://serverfault.com/questions/490261/cannot-use-any-yum-command-no-space-left-on-device>


upvoted 3 times

Which of the following is an example of load balancing?

- A. Round robin
- B. Active-active
- C. Active-passive
- D. Failover

Suggested Answer: B

Community vote distribution



proservices 4 months ago

Selected Answer: A

ChatGPT

The correct answer is:

A. Round robin

Explanation:

Load balancing is the process of distributing network traffic or workloads across multiple servers to optimize resource use, improve performance, and ensure high availability.

Round robin (A) is a load balancing technique where requests are distributed sequentially across a group of servers. It helps balance the workload among multiple servers.

Active-active (B) refers to a high availability configuration where multiple servers actively handle traffic, which can include load balancing but is not a specific technique itself.

Active-passive (C) refers to a setup where one server is active while another remains on standby, which is more related to redundancy than load balancing.

Failover (D) is a redundancy mechanism where a backup system takes over when the primary system fails, not a load balancing method.

Thus, Round robin is the best example of load balancing.

upvoted 1 times

Stache 6 months ago

Selected Answer: A

(A) Round Robin is the only Load Balancing example, the rest are all Server Cluster terms.

upvoted 1 times

fluke92 7 months, 1 week ago

Selected Answer: A

Load balancing is the process of distributing network or application traffic across multiple servers to ensure no single server is overwhelmed. Round robin is a specific load balancing method where incoming requests are distributed sequentially across available servers in a cyclical manner.

Why the Other Options Are Incorrect:

B. Active-active:

This is a high-availability configuration where multiple servers or nodes are active simultaneously, but it focuses on redundancy and availability rather than explicitly on distributing traffic.

C. Active-passive:

This refers to a failover configuration where one server is active while the other is passive and only takes over in case of failure. This is not a load balancing technique.

D. Failover:

Failover is a method to switch to a backup system when the primary system fails. It ensures availability but does not distribute traffic among multiple servers.

upvoted 1 times

🗳️ 👤 **kx7tg4xu** 10 months, 2 weeks ago

Selected Answer: B

B is the correct answer to this question

upvoted 1 times

🗳️ 👤 **Ckamunga** 1 year, 11 months ago

The correct answer is A

upvoted 1 times

🗳️ 👤 **K1lroy** 2 years, 1 month ago

Selected Answer: A

They are asking for an example to LB. Active-active should be load balanced of course but it's not an example..it's basically a requirement while "round robin" is an example of multiple choices that are available

upvoted 1 times

🗳️ 👤 **AbusedInk** 2 years, 4 months ago

A

upvoted 2 times

🗳️ 👤 **Obi_Wan_Jacoby** 2 years, 5 months ago

Selected Answer: B

I agree with King2 (Answer B is it). Answer A (Round Robin) is an algorithm. While Answer B is the top solution. The algorithm is only able to work if the active-active cluster is there, which is the load balancing solution.

upvoted 1 times

🗳️ 👤 **tem_knows** 1 year, 3 months ago

Active-active and Active-passive are configurations for high availability and redundancy rather than load balancing. They involve using multiple servers where all (active-active) or some (active-passive) are ready to take over in case of a failure but don't inherently distribute load for efficiency.

upvoted 1 times

🗳️ 👤 **Pongsathorn** 2 years, 7 months ago

Selected Answer: A

BALANCING

There are many scheduling types for load balancers. Three of them are round robin, most recently used, and weighted scheduling.

- Round robin
- Most recently used (MRU)
- Weighted scheduling

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 244.

upvoted 2 times

🗳️ 👤 **Pongsathorn** 2 years, 7 months ago

(cont.)

TYPES OF SERVER CLUSTERS

There are many server cluster variations. Common differences include the number of participating nodes, whether or not a node is active during normal operations, and how a cluster determines normal activities.

There are two configuration types for clusters: active-active and active-passive.

- Active-active clusters use a load balancer device to distribute server requests between the nodes. This design provides high availability and also increases performance. The load balancer distributes the workload among the cluster nodes.

- Active-passive clusters provide only high availability and not load balancing. All requests go to the active server. If that node fails, however, the requests go to the second server.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 241.

The question is asking about "load balancing" not "server cluster".

upvoted 1 times

🗨️ 👤 **King2** 2 years, 9 months ago

Selected Answer: B

Round robin refers to a load balancer schedule, which assigns connections in order.

Weighted scheduling also refers to a load balancer schedule and is used when the server's hardware capabilities vary, allocating more connections to the more powerful servers.

upvoted 3 times

🗨️ 👤 **jagoichi** 2 years, 9 months ago

Active-Active is correct - Round Robin is a scheduler

Active-active clusters use a load balancer device to distribute server requests between the nodes. This design provides high availability and also increases performance. The load balancer distributes the workload among the cluster nodes.

There are many scheduling types for load balancers. Three of them are round robin, most recently used, and weighted scheduling

upvoted 3 times

Which of the following is the MOST appropriate scripting language to use for a logon script for a Linux box?

- A. VBS
- B. Shell
- C. Java
- D. PowerShell
- E. Batch

Suggested Answer: B

Community vote distribution

B (100%)

 **Slappy1Eye** Highly Voted 3 years, 5 months ago

The question is: Which of the following is the MOST appropriate scripting language to use for a logon script for a Linux box? The answer shows A: VBS

The link below goes to a site about Active Directory information.

I believe the answer is wrong and should be Shell since VBS is a Microsoft based script.

Can this answer be verified?

upvoted 11 times

 **dnc1981** Highly Voted 3 years, 3 months ago

Shell is definitely the correct answer

upvoted 8 times

 **Pongsathorn** Most Recent 9 months ago

Selected Answer: B

Login and Logout Scripts

Some scripts execute based on the user. For example, a login script runs when a user logs in. The script customizes the user's environment, maps network drives, and automatically launches a particular application. Logout scripts execute when the user signs off. These scripts might delete temp files or copy the user's data to a network server.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 160.

SHELL LANGUAGES

Scripts are written using a shell language. The shell is the server's command-line environment. The default shell for Linux is bash. Windows has two shells: cmd.exe and PowerShell, although Microsoft emphasizes PowerShell. As a general rule, a script written for one shell, such as PowerShell, will not be understood by another shell, such as bash. The system must have the appropriate shell installed to process the script.

upvoted 2 times

 **Pongsathorn** 2 years, 7 months ago

(cont.)

Types of scripts:

Batch file -- a series of commands, usually based on MS-DOS, that execute with limited additional functionality.

Bash script -- a script written in bash language that contains commands, variables, constructs, and other components. Bash scripts are almost always written for Linux systems.

PowerShell script -- a script written in the PowerShell language that contains cmdlets (PowerShell commands), variables, constructs, and other components. PowerShell scripts are almost always written for Windows systems.

VBS script -- a script written in the VBScript language. Microsoft has deprecated VBScript, but there have been many VBScript scripts written over the years. VBScript scripts are almost always written for Windows systems.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 154.

upvoted 2 times

🗳️ 👤 **kx7tg4xu** 10 months, 2 weeks ago

Selected Answer: B

B is the correct answer to this question
upvoted 1 times

🗳️ 👤 **Sweety_Certified7** 1 year, 5 months ago

Selected Answer: B

Shell script can be used for logon script for Linux
upvoted 1 times

🗳️ 👤 **AbusedInk** 2 years, 4 months ago

A.
Shell is not a scripting language.
upvoted 2 times

🗳️ 👤 **Sweety_Certified7** 1 year, 5 months ago
Shell script can be used for logon script for Linux
upvoted 1 times

🗳️ 👤 **lordguck** 2 years, 7 months ago
It's B. D is more DOS/Windows like (.bat/.cmd)
upvoted 1 times

🗳️ 👤 **lordguck** 2 years, 7 months ago
sorry, DE is more DOS/windows like (.bat/cmd/powershell),
upvoted 1 times

🗳️ 👤 **nixonbii** 2 years, 8 months ago

Selected Answer: B

Just look up the uses of VBS for yourself. BASH CLI is to Linux what peanut butter is to jelly.
upvoted 1 times

🗳️ 👤 **paperburn** 2 years, 8 months ago

Selected Answer: B

shell is correct
upvoted 1 times

🗳️ 👤 **szl0144** 2 years, 9 months ago

Selected Answer: B

must be shell
upvoted 1 times

🗳️ 👤 **Rainkoot** 2 years, 10 months ago

Selected Answer: B

I would go with Linux shell.
upvoted 1 times

🗳️ 👤 **Dion79** 3 years, 2 months ago

Selected Answer: B

I would go with Linux shell.
upvoted 2 times

🗳️ 👤 **Atemius** 3 years, 3 months ago

I agree the answer should be B- Shell.
upvoted 3 times

🗳️ 👤 **Ariel235788** 3 years, 3 months ago

Definitely Shell. VBS, batch, and powershell are Microsoft based
upvoted 3 times

🗳️ 👤 **Dion79** 3 years, 4 months ago

I would go with B. The Answer is incorrect. I would go with Shell or Bash is always used for Linus.

Scripts are written using a shell language. The shell is the server's command-line environment. The default shell for Linux is bash. Windows has two shells: cmd.exe and PowerShell, although Microsoft emphasizes PowerShell. As a general rule, a script written for one shell, such as PowerShell, will not be understood by another shell, such as bash. The system must have the appropriate shell installed to process the script.

Bash script -- a script written in bash language that contains commands, variables, constructs, and other components. Bash scripts are almost always written for Linux systems.

VBS script -- a script written in the VBScript language. Microsoft has deprecated VBScript, but there have been many VBScript scripts written over the years. VBScript scripts are almost always written for Windows systems.

upvoted 2 times

  **maffo02** 3 years, 5 months ago

She'll should be the correct answer.

upvoted 2 times

Which of the following tools will analyze network logs in real time to report on suspicious log events?

- A. Syslog
- B. DLP
- C. SIEM
- D. HIPS

Suggested Answer: C

Reference:

<https://www.manageengine.com/products/eventlog/syslog-server.html>

Community vote distribution

C (100%)

🗳️ 👤 **Grumpy_Old_Coot** 9 months ago

Learn the acronyms and the answer just falls out. You don't even need to know what SIEM means.

Syslog - Stores a log file.

HIPS - Host Intrusion Prevention System

DLP - Data Loss Prevention

upvoted 1 times

🗳️ 👤 **Ckamunga** 11 months, 1 week ago

The correct answer is C

upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 1 year, 5 months ago

Selected Answer: C

I concur with C

upvoted 1 times

🗳️ 👤 **Atemius** 2 years, 3 months ago

A-SIEM is the correct answer

"Security information and event management (SIEM) software provides a centralized repository for logs, audit events, and security device alerts to detect and notify admins of suspicious activity. For example, attackers can enable a backdoor on a compromised device that will provide them with undetected access for long periods of time; this problem is often solved by applying patches, but SIEM solutions may detect repeated abnormal login times to a server."

upvoted 2 times

🗳️ 👤 **Atemius** 2 years, 3 months ago

Sorry, meant C is the correct answer. SIEM logs

upvoted 1 times

🗳️ 👤 **JRod42** 2 years, 3 months ago

Answer says SIEM but the link goes to SYSLOG?

upvoted 1 times

🗳️ 👤 **Ariel235788** 2 years, 3 months ago

Syslog logs can be monitored by the SIEM. SIEM would do the alerting. Syslog is just raw data. The 2nd part of the question says to alert on sus activity

upvoted 1 times

🗳️ 👤 **Dion79** 2 years, 4 months ago

Linux log files are generated by a service named "rsyslog." By default, rsyslog stores log files in the /var/log directory. That directory contains many log files, though its exact contents vary depending on the Linux distribution and the installed applications.

SIEM - Hardware failure Log analysis (SIEM), redundancy. Malware Data monitoring, log analysis (SIEM). Data corruption Data monitoring, backups. Insider threats Log analysis (SIEM), two-person integrity, separation of roles, role rotation, regulatory constraints. Data Loss Prevention (DLP). Data

monitoring, two-person integrity, regulatory constraints, data retention. Unwanted duplication (theft) Log analysis (SIEM), data monitoring, two-person integrity, regulatory constraints, data retention. Unwanted publication (theft) Log analysis (SIEM), data monitoring, two-person integrity, regulatory constraints, data retention. Unwanted access (backdoor, social engineering). Log analysis (SIEM), two-person integrity Breach Log analysis (SIEM), data monitoring, regulatory constraints, data retention.

upvoted 2 times

Which of the following will correctly map a script to a home directory for a user based on username?

- A. \\server\users\$\username
- B. \\server%\username%
- C. \\server\FirstInitialLastName
- D. \\server\\$\username\$

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ **Kraken84** 8 months, 3 weeks ago

The correct option that will map a script to a home directory for a user based on the username is:

B. \\server%\username%

In this option, `%username%` is a variable that will be replaced with the actual username of the user. This allows for dynamic mapping based on the logged-in user's username.

~ChatGPT 4.0

upvoted 3 times

🗨️ **PEsty93** 2 years, 2 months ago

Selected Answer: B

% are used for variables in Microsoft. \$ are Linux/Unix

upvoted 4 times

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

Yeah this question is horrid, because syntax is dependant upon the type of shell you are using/language. so multiple possible answers? I thought \$ is for Linux as well. This is one of the worst CompTIA questions I've ever seen.

upvoted 3 times

🗨️ **Atemius** 2 years, 3 months ago

Correct answer is B. The link shows the answer.

upvoted 3 times

A server that recently received hardware upgrades has begun to experience random BSOD conditions. Which of the following are likely causes of the issue?

(Choose two.)

- A. Faulty memory
- B. Data partition error
- C. Incorrectly seated memory
- D. Incompatible disk speed
- E. Uninitialized disk
- F. Overallocated memory

Suggested Answer: AC

Community vote distribution

AC (100%)

🗲️ 👤 **Laudy** 9 months, 4 weeks ago

Selected Answer: AC

AC was my answer too

upvoted 1 times

🗲️ 👤 **Laudy** 9 months, 4 weeks ago

AC was my answer too

upvoted 1 times

🗲️ 👤 **Jus2Lewis** 1 year ago

A & C are the correct answers.

upvoted 2 times

A server administrator has configured a web server. Which of the following does the administrator need to install to make the website trusted?

- A. PKI
- B. SSL
- C. LDAP
- D. DNS

Suggested Answer: B

Community vote distribution

B (60%)

A (40%)

🗳️ 👤 **[Removed]** 10 months, 1 week ago

Selected Answer: B

The correct answer is B.

Because the fact that SSL/TLS is only one little invention based on PKI, A is not the exactly precise answer here.

upvoted 1 times

🗳️ 👤 **kx7tg4xu** 10 months, 2 weeks ago

Selected Answer: B

The correct answer is B

upvoted 1 times

🗳️ 👤 **RBL23168** 1 year, 2 months ago

Selected Answer: B

B. SSL

You don't INSTALL PKI. Security in PKI is done with SSL certificates. SSL (Secure Sockets Layer) is the security protocol used on the web when you fetch a page whose address begins with https.

upvoted 1 times

🗳️ 👤 **Grumpy_Old_Coot** 1 year, 9 months ago

PKI = Public Key Infrastructure (Certificates)

SSL = Secure Socket Layer (uses PKI certificates) to secure access by identifying a specific server/workstation/user.

Look -real- carefully at what this question is asking...

upvoted 1 times

🗳️ 👤 **Dingos** 2 years, 1 month ago

Selected Answer: A

I was going also for B, BUT than I realised that with PKI you can issue SSL.

Now I am 65% sure than ansver is PKI as with it you can get both.

upvoted 1 times

🗳️ 👤 **Dingos** 2 years, 1 month ago

and minute later I was back on B as that red question and decided that SSL is right answer as after SSL install website becoming trusted.

upvoted 2 times

🗳️ 👤 **gingasaurusrex** 2 years, 2 months ago

ChatGPT, when asked this exact question, went with SSL

upvoted 2 times

🗳️ 👤 **Laudy** 2 years, 3 months ago

Selected Answer: A

[https://www.thesslstore.com/blog/what-is-pki-a-crash-course-on-public-key-infrastructure-](https://www.thesslstore.com/blog/what-is-pki-a-crash-course-on-public-key-infrastructure-pki/#:~:text=These%20certificates%20are%20known%20as,encryption%20encrypts%20the%20message%20itself.)

[pki/#:~:text=These%20certificates%20are%20known%20as,encryption%20encrypts%20the%20message%20itself.](https://www.thesslstore.com/blog/what-is-pki-a-crash-course-on-public-key-infrastructure-pki/#:~:text=These%20certificates%20are%20known%20as,encryption%20encrypts%20the%20message%20itself.)

When answering the question "what is PKI?" you need to talk not only about what it constitutes but also how it's used. There are several ways that businesses and organizations around the world use public key infrastructure:

There are several different types of digital certificates, which (you may notice correspond to the popular uses of PKI that we covered above). Here are some of the types of X.509 digital certificates that you can find within the PKI infrastructure:

SSL/TLS certificates
S/MIME certificates
Code signing certificates
Client authentication certificates

In short, PKI includes much more than SSL, and is the more correct answer.

upvoted 1 times

🗲️ 👤 **gingasaurusrex** 2 years, 4 months ago

Selected Answer: A

SSL cert

upvoted 1 times

🗲️ 👤 **Obi_Wan_Jacoby** 2 years, 5 months ago

Selected Answer: B

Answer is B. (tricky question) You must purchase/install a SSL cert. Along with it comes the PKI keys (Public Key Infrastructure). The PKI allows the message to be encrypted, but the SSL cert PROVES the site is who they claim to be (A.K.A TRUSTED..which the questions states is needed).

<https://www.thesslstore.com/blog/what-is-pki-a-crash-course-on-public-key-infrastructure-pki/>

upvoted 1 times

🗲️ 👤 **Pongsathorn** 2 years, 7 months ago

Selected Answer: B

Web interface—this may also manage Linux servers. For example, Linux Cockpit allows you to display and manage your Linux servers via a web browser. Such administration is therefore relatively universal, since the administrator only has to be working from a device that can run a web browser. SSL VPN protocols may be used to secure the connection.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 199.

CompTIA Server+ Certification Exam Objectives EXAM NUMBER: SK0-005 never mentioned PKI, so the best answer is SSL.

upvoted 2 times

🗲️ 👤 **paperburn** 2 years, 8 months ago

Selected Answer: A

A PKI certificate is a trusted digital identity. It is used to identify users, servers or things when communicating over untrusted networks, to sign code or documents and to encrypt data or communication. A PKI certificate is also called a digital certificate.

upvoted 1 times

🗲️ 👤 **jagoichi** 2 years, 9 months ago

I agree Timock. This is referencing Certificates

upvoted 1 times

🗲️ 👤 **Timock** 2 years, 9 months ago

You don't INSTALL PKI. Security in PKI is done with SSL certificates. SSL (Secure Sockets Layer) is the security protocol used on the web when you fetch a page whose address begins with https:. TLS (Transport Layer Security) is a newer version of the protocol. In practice, most websites now use the new version.

upvoted 2 times

A technician is attempting to update a server's firmware. After inserting the media for the firmware and restarting the server, the machine starts normally into the

OS. Which of the following should the technician do NEXT to install the firmware?

- A. Press F8 to enter safe mode
- B. Boot from the media
- C. Enable HIDS on the server
- D. Log in with an administrative account

Suggested Answer: B

 **Atemius**  9 months, 1 week ago

B-Boot from Media is the correct answer

Firmware updates (such as BIOS or the Unified Extensible Firmware Interface [UEFI]) may be required prior to installation, even for indirect reasons, such as network Preboot Execution Environment (PXE) boot, USB device boot support, and so on. If you plan on installing the hypervisor from a DVD or USB, you may also have to change the boot order configuration on your machine.

upvoted 7 times

A server administrator mounted a new hard disk on a Linux system with a mount point of /newdisk. It was later determined that users were unable to create directories or files on the new mount point. Which of the following commands would successfully mount the drive with the required parameters?

- A. echo /newdisk >> /etc/fstab
- B. net use /newdisk
- C. mount -o remount, rw /newdisk
- D. mount -a

Suggested Answer: C

Reference:

<https://unix.stackexchange.com/questions/149399/how-to-remount-as-read-write-a-specific-mount-of-device>

Community vote distribution

 C (100%)

🗳️ 👤 **fluke92** 7 months, 1 week ago

Selected Answer: C`mount -o remount,rw /newdisk`

upvoted 1 times

🗳️ 👤 **tame_rabbit** 1 year, 3 months ago

The correct command to mount the drive with the required parameters is:

C. `mount -o remount,rw /newdisk`

Explanation:

`mount`: This command is used to mount filesystems in Linux.`-o remount,rw`: This option remounts the filesystem in read-write mode. It's necessary here because the issue seems to be related to users being unable to create directories or files, which suggests that the disk might have been mounted as read-only.`/newdisk`: This is the mount point where the new disk is mounted.By using the `-o remount,rw` options, we ensure that the disk is remounted with read-write permissions, allowing users to create directories and files on the mount point `/newdisk`

upvoted 1 times

🗳️ 👤 **dddc1000** 2 years, 9 months ago

My bad, should be lowercase letter o

`mount -o remount,rw /newdisk`

upvoted 3 times

🗳️ 👤 **dddc1000** 2 years, 9 months ago

Answer is c. The correct syntax is:

`mount -o remount,rw /newdisk`

upvoted 1 times

🗳️ 👤 **King2** 2 years, 9 months ago

A. `echo /newdisk >> /etc/fstab`B. `net use /newdisk`C. `mount -o remount, rw /newdisk`D. `mount -a`

upvoted 2 times

🗳️ 👤 **szl0144** 2 years, 9 months ago

Selected Answer: C

I will choose C

upvoted 1 times

🗳️ 👤 **Betny77** 3 years, 3 months ago

I found no evidence of the symbol in the given answer in the man page for the mount command. I think the answer is A. Anyone else?
upvoted 1 times

🗨️ 👤 **Ariel235788** 3 years, 3 months ago

the symbol is a '-' for whatever reason when they upload answer choices it gets jacked up to the weird symbols.
upvoted 3 times

🗨️ 👤 **PEsty93** 3 years, 3 months ago

"Echo" wouldn't perform any command?
upvoted 3 times

🗨️ 👤 **dnc1981** 3 years, 3 months ago

C is the correct answer. The funny character is supposed to be a hyphen
upvoted 3 times

Which of the following BEST describes the concept of right to downgrade?

- A. It allows for the return of a new OS license if the newer OS is not compatible with the currently installed software and is returning to the previously used OS
- B. It allows a server to run on fewer resources than what is outlined in the minimum requirements document without purchasing a license
- C. It allows for a previous version of an OS to be deployed in a test environment for each current license that is purchased
- D. It allows a previous version of an OS to be installed and covered by the same license as the newer version

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **haazybanj** Highly Voted 2 years, 9 months ago

D is correct

upvoted 9 times

🗳️ 👤 **Pieman125** Highly Voted 3 years, 2 months ago

I think the correct Answer is D

upvoted 7 times

🗳️ 👤 **warmlazana** Most Recent 8 months, 4 weeks ago

Selected Answer: D

D is actual definition

upvoted 1 times

🗳️ 👤 **cloudchief25** 1 year ago

Sadly B is correct. On the real exam the correct answer is B.

While yes B is a true statement. The question is what BEST describes it. And in reality, the best answer is D

upvoted 2 times

🗳️ 👤 **iTomi** 1 year, 2 months ago

Selected Answer: D

What are Downgrade Rights?

Downgrade Rights are where end users who have acquired the Latest Version of some kind of software can use an earlier version of that same software, until they are ready to migrate to the later version of the Software/Technology.

upvoted 1 times

🗳️ 👤 **jjwelch00** 1 year, 7 months ago

D is correct

upvoted 1 times

🗳️ 👤 **kloug** 1 year, 9 months ago

dddddddddd

upvoted 1 times

🗳️ 👤 **Mel152** 1 year, 9 months ago

If D is the correct answer why are they stating B. Why don't they update it to be B this is very confusing !!!

upvoted 2 times

🗳️ 👤 **gingasaurusrex** 1 year, 10 months ago

Selected Answer: D

D is right here

upvoted 1 times

🗳️ 👤 **paperburn** 2 years, 2 months ago

Selected Answer: D

If a product includes downgrade rights, the license terms for that product will indicate which earlier versions of the software may be used.

upvoted 1 times

🗨️ 👤 **Noms100** 2 years, 6 months ago

D is correct, If a product includes downgrade rights, the license terms for that product will indicate which earlier versions of the software may be used.

upvoted 2 times

🗨️ 👤 **[Removed]** 2 years, 8 months ago

Selected Answer: D

Definitely D

upvoted 2 times

🗨️ 👤 **dnc1981** 2 years, 9 months ago

Definitely D is the correct answer

upvoted 3 times

A server administrator needs to harden a server by only allowing secure traffic and DNS inquiries. A port scan reports the following ports are open:

443

636

Which of the following open ports should be closed to secure the server properly? (Choose two.)

A. 21

B. 22

C. 23

D. 53


E. 443

F. 636

Suggested Answer: AC

Community vote distribution

AC (100%)

 **Pongsathorn** Highly Voted 2 years, 7 months ago

Selected Answer: AC

21 - FTP

22 - SSH

23 - Telnet


53 - DNS

443 - HTTPS

636 - LDAPS

There are 3 ports which not secure and need to close 21, 23, 53 but the server provide DNS inquiries, so we don't close port 53.

upvoted 6 times

 **Musa007** Most Recent 2 years, 1 month ago

DNS can also use other ports for specialized purposes. For example, DNS over TLS (DoT) uses port 853, and DNS over HTTPS (DoH) typically uses port 443 to encapsulate DNS traffic within HTTPS


That's why 53 is blocked

upvoted 1 times

 **ccoli** 11 months, 1 week ago

They shouldn't have ftp in the answer results. But since they said securing traffic and DNS I think the puposed answer of CD and is what they're looking for, they're just made it a trick question to try to glean more exam fees from people for a cert no one respects.

upvoted 1 times

 **Jfrican** 2 years, 3 months ago

Then get a new scanner because it is not showing Ports 21 and 23

upvoted 2 times

 **zozo1978** 2 years, 4 months ago

How could DNS Port 53 needs to be close where the question asking for DNS Port to be used ?

upvoted 1 times

 **nixonbii** 2 years, 8 months ago

Selected Answer: AC

A - Why would you leave an unsecure FTP port open? C - Telnet is notorious for sending credentials over the network in plain text. If you close port 53 how will hosts on the network get name resolution?

upvoted 1 times

 **paperburn** 2 years, 8 months ago

Selected Answer: AC

No ftp or telnet allowed

upvoted 1 times

🗨️ 👤 **Fineb** 2 years, 9 months ago

A and C should be the best answer

upvoted 1 times

🗨️ 👤 **jagoichi** 2 years, 9 months ago

AC

Port 21 and File Transfer

FTP is often thought of as a “not secure” file transfer protocol. This is mainly due to FTP sending data in clear text and offering an anonymous option with no password required. However, FTP is a trusted and still widely used protocol for transferring files

Is port 23 secure?

Port 23 – Telnet. A predecessor to SSH, is no longer considered secure and is frequently abused by malware.

upvoted 1 times

🗨️ 👤 **szl0144** 2 years, 9 months ago

53 is for DNS query, why we need to close it?

upvoted 1 times

A server administrator wants to check the open ports on a server. Which of the following commands should the administrator use to complete the task?

- A. nslookup
- B. nbtstat
- C. telnet
- D. netstat -a

Suggested Answer: D

Reference:

<https://linuxhint.com/netstat-a/>

Community vote distribution

D (100%)

🗳️ 👤 **San24Yeah** 7 months, 2 weeks ago

D is the correct answer. netstat -a will return all the active ports
upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 1 year, 11 months ago

Selected Answer: D

D is correct
upvoted 1 times

Which of the following must a server administrator do to ensure data on the SAN is not compromised if it is leaked?

- A. Encrypt the data that is leaving the SAN
- B. Encrypt the data at rest
- C. Encrypt the host servers
- D. Encrypt all the network traffic

Suggested Answer: B

Community vote distribution

B (67%)

A (33%)

🗳️ 👤 **fluke92** 7 months, 1 week ago

Selected Answer: B

Encrypting data at rest ensures that the data stored on the SAN is protected from unauthorized access, even if the storage medium is compromised or stolen. This is a critical step to secure sensitive information and mitigate the risk of data breaches.

upvoted 2 times

🗳️ 👤 **a792193** 1 year, 6 months ago

Selected Answer: B

B. Encrypting Data at rest would ensure the data is unreadable if it is leaked.

upvoted 1 times

🗳️ 👤 **Kraken84** 1 year, 8 months ago

To ensure data on the SAN (Storage Area Network) is not compromised if it is leaked, the server administrator should:

B. Encrypt the data at rest

Encrypting the data at rest ensures that even if the data is accessed or leaked, it remains unreadable without the appropriate decryption key. This provides a layer of security against unauthorized access to the data stored on the SAN.

~ChatGpT4

upvoted 1 times

🗳️ 👤 **Amorprecious** 2 years, 3 months ago

B. To ensure data on the SAN is not compromised if it is leaked, a server administrator should encrypt the data at rest. This means that the data is encrypted while it is stored on the SAN. This way, even if the data is accessed or copied without authorization, it will be unreadable without the encryption key.

Encrypting the data that is leaving the SAN or encrypting all network traffic can help protect data while it is in transit, but it does not necessarily protect the data if it is leaked from the SAN. Encrypting the host servers may also be a good security practice, but it does not directly address the issue of data leakage from the SAN.

upvoted 1 times

🗳️ 👤 **Laudy** 2 years, 3 months ago

Selected Answer: B

It's referring to data ON the SAN.

Data at rest.

upvoted 2 times

🗳️ 👤 **Obi_Wan_Jacoby** 2 years, 5 months ago

Selected Answer: B

I concur with answer B. I see answer D achieving the need for SAN traffic being encrypted. Why have A and D as D would take care of said traffic? The encryption of data at rest is what we are looking for.

upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 2 years, 5 months ago

Coming back to this a little later.. Now I am thinking answer A is in-fact correct.

upvoted 1 times

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

I can accept B but I wish they had not included D as an option. Basic cyber-security principle: protection of data in motion, at rest, in process.

upvoted 2 times

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

Encrypt Data at Rest on the Storage Array of drives.

upvoted 2 times

🗨️ 👤 **dcdc1000** 2 years, 9 months ago

Answer is B. Because the question states "data on the SAN" so if you encrypt data at rest, this will prevent data compromised.

upvoted 2 times

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

Selected Answer: A

Best Guess is A, Encrypt all data leaving the SAN. Data can leave through local transfer (USB drive) a network vulnerability from outsider/insider, . So data needs to be encrypted at rest, but also possibly in transit. But if you can encrypt "all data that leaves SAN" then that would cover everything... all these questions seem very ambiguous.....

upvoted 2 times

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

SAN is a network, the whole idea of this approach is for multiple devices to be able to access a shared filesystem, not DAS,, which is the Storage device attached to a server that SAN can use to share a filesystem across multiple devices on the network

Data Loss(breach) Prevention Method:

1; Encrypt all "SAN" network traffic..

2; Another approach is to Encrypt Data at Rest on the Storage Array of drives.

They both go together like peanut butter and jelly for a perfect sandwich (Data Breach Protection on SAN).

upvoted 1 times

A server technician has been asked to upload a few files from the internal web server to the internal FTP server. The technician logs in to the web server using PuTTY, but the connection to the FTP server fails. However, the FTP connection from the technician's workstation is successful. To troubleshoot the issue, the technician executes the following command on both the web server and the workstation: `ping ftp.acme.local`

The IP address in the command output is different on each machine. Which of the following is the MOST likely reason for the connection failure?

- A. A misconfigured firewall
- B. A misconfigured hosts.deny file
- C. A misconfigured hosts file
- D. A misconfigured hosts.allow file

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Grumpy_Old_Coot** 9 months ago

Selected Answer: C

Hosts.Allow has nothing to do with "ping" command. The host file is a "cruddy DNS" on a local machine (both Windows and Linux). Different Host files = different DNS lookup results.

upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 1 year, 5 months ago

Selected Answer: C

C is correct

upvoted 1 times

🗳️ 👤 **sotodaniel77** 1 year, 8 months ago

Selected Answer: C

C is the answer

upvoted 1 times

🗳️ 👤 **szl0144** 1 year, 9 months ago

Selected Answer: C

C is the answer

upvoted 1 times

🗳️ 👤 **Rainkoot** 1 year, 10 months ago

Selected Answer: C

Remove PuTTY from the question it's intended as a distraction.

The key is "ping ftp.acme.local"

The IP address in the command output is different on each machine."

hosts.allow and hosts.deny are for ACL not direction

a host file that is different on each machine would cause the IP address to be different

upvoted 3 times

🗳️ 👤 **dnc1981** 2 years, 3 months ago

C is the answer

upvoted 3 times

🗳️ 👤 **PEsty93** 2 years, 2 months ago

No, the answer is D. hosts file is used in Windows. host.allow is used in Linux.



if you're connecting the server using Putty that means the server is Linux.

upvoted 4 times

🗳️ 👤 **Replicant** 1 year, 3 months ago

Wrong on the Putty / linux statement. It isn't used to connect ssh only to a linux box. It can be used ssh to a windows server as well.

upvoted 1 times



  **i_bird** 1 year, 10 months ago

Windows 10 Host File is found at c:\Windows\System32\Drivers\etc\hosts

Linux host file is found at /etc/hosts..

C: Misconfigured host file

upvoted 1 times

  **Dion79** 2 years ago

C might be the answer. No reference of Linux, and Putty also can be used for Windows & Linux.

upvoted 2 times

A company deploys antivirus, anti-malware, and firewalls that can be assumed to be functioning properly. Which of the following is the MOST likely system vulnerability?

- A. Insider threat
- B. Worms
- C. Ransomware
- D. Open ports
- E. Two-person integrity

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **Grumpy_Old_Coot** 9 months ago

Selected Answer: A

Users (insider threat) are the biggest problem on any properly configured network.

upvoted 1 times

🗳️ 👤 **King2** 1 year, 9 months ago

Selected Answer: A

Correct answer: A

The company implemented technical controls (antivirus, Anti-malware, and firewalls) and these controls are working properly.

The remaining vulnerability would be insider threat. Internal user threats (careless insider)

Users who unknowingly expose the system to outside threats. This is the most common type of insider threat, resulting from mistakes, such as leaving a device exposed or falling victim to a scam. For example, an employee who intends no harm may click on an insecure link, infecting the system with malware.

The biggest security vulnerability in any organization is its own employees. Whether it's the result of intentional malfeasance or an accident, most data breaches can be traced back to a person within the organization that was breached.

For example, employees may abuse their access privileges for personal gain. Or, an employee may click on the wrong link in an email, download the wrong file from an online site, or give the wrong person their user account credentials—allowing attackers easy access to your systems.

upvoted 2 times

🗳️ 👤 **Timock** 1 year, 9 months ago

Selected Answer: A

The rest of the options besides Insider Threats are covered by the "...can be assumed to be functioning properly."

upvoted 1 times

🗳️ 👤 **i_bird** 1 year, 10 months ago

keywords are important in CompTIA exams, They are identified by the capitalization of the keyword, in this case, it is MOST likely system vulnerabilities.

Inside Threats seem like the MOST likely system vulnerabilities, Cause Open ports vulnerabilities have been partly mitigated using the anti-malware program, the other possible attack using open ports is social engineering attacks which can basically be mitigated using awareness programs for the users.



There is no inside threat attack mitigation process outlined in the question, so I think inside threat is the MOST likely system vulnerability from the options outlined.

upvoted 1 times

🗳️ 👤 **PEsty93** 2 years, 3 months ago

The wording on this question leaves a lot open for assumption. Just because you have a Firewall doesn't mean you don't have open ports

upvoted 2 times

  **Rainkoot** 1 year, 10 months ago

It does leave it open for assumption however, with "...assumed to be functioning properly."

I also believe them to be implying configured correctly with bare minimum needed for business operations.

upvoted 1 times

A security analyst suspects a remote server is running vulnerable network applications. The analyst does not have administrative credentials for the server. Which of the following would MOST likely help the analyst determine if the applications are running?

- A. User account control
- B. Anti-malware
- C. A sniffer
- D. A port scanner

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ **dnc1981** Highly Voted 2 years, 9 months ago

D is the correct answer
upvoted 10 times

🗳️ **a792193** Most Recent 1 year ago

Selected Answer: D

D. Port Scanner. You can run nmap to determine what ports are open and what services are running on them.
upvoted 1 times

🗳️ **Pongsathorn** 2 years, 1 month ago

Selected Answer: D

User Account Control (UAC)

Administrator and root accounts are the most highly privileged accounts in an operating system. When a server is left logged on with a privileged account, it creates a huge security issue. Most of the server operating systems you will encounter today incorporate the ability of an administrator or a root account holder to use a nonprivileged account as standard operating procedure and elevate their privileges as needed without logging off and logging back in as root.

The User Account Control feature in Windows and the use of the sudo command in Linux make this possible. Using either system an administrator can elevate their privileges for a specific task and that security context ends when they are finished with that task.

In Windows this can be done in the GUI by right-clicking the icon representing the task and selecting Run As Administrator

UAC REQUIRES administrator-level access, provided answer is not correct.
upvoted 1 times

🗳️ **Pongsathorn** 2 years, 1 month ago

Selected Answer: D

Internet Control Message Protocol (ICMP) messages can be used to scan a network for open ports. Open ports indicate services that may be running and listening on a device that may be susceptible to attack. An ICMP, or port scanning, attack basically pings every address and port number combination and keeps track of which ports are open on each device as the pings are answered by open ports with listening services and not answered by closed ports. One of the most widely used port scanners is Network Mapper (Nmap), a free and open source utility for network discovery and security auditing.

upvoted 2 times

🗳️ **nixonbii** 2 years, 2 months ago

Never seen UAC as a way to detect the status of any network application.
upvoted 1 times

🗳️ **Dion79** 2 years, 7 months ago

Provided answer may be correct. If the analyst does not have administrative credentials for the server, will he/she be able to run the port scan on the Server/Network without admin credentials?

You can look at the UAC to see if applications are running.

Reference:

<https://docs.microsoft.com/en-us/windows/security/identity-protection/user-account-control/how-user-account-control-works>

upvoted 2 times

  **TheITStudent** 2 years, 4 months ago

@Dion... close, but not correct... you were on the right track. UAC or user account control is basically Mandatory Access Control. which REQUIRES administrator level access. That would be the correct answer IF "analyst does not have administrative credentials" was not in there. So the analyst needs to verify/check vulnerable applications without administrative access. welcome to the world of hackers... you run nmap or something of the sort. sniffing would work, but if traffic is encrypted, it might be more challenging. best answer is to do a port scan. D is my answer

upvoted 5 times

  **Ariel235788** 2 years, 9 months ago

Fairly certain a port scanner would be used here. You can run nmap to determine what ports are open and what services are running on them. could even run a -sV to determine version

upvoted 4 times

  **Ariel235788** 2 years, 9 months ago

UAC would only prevent installs. As stated in the reference link. User Account Control: Detect application installations and prompt for elevation. Does not check if an application is already installed.

upvoted 1 times

  **Ariel235788** 2 years, 9 months ago

'MOST likely help the analyst determine if the applications are running?' not being run. Answer choice D makes the most logical sense in terms of the question context

upvoted 2 times

A server is performing slowly, and users are reporting issues connecting to the application on that server. Upon investigation, the server administrator notices several unauthorized services running on that server that are successfully communicating to an external site. Which of the following are MOST likely causing the issue? (Choose two.)

- A. Adware is installed on the users' devices
- B. The firewall rule for the server is misconfigured
- C. The server is infected with a virus
- D. Intrusion detection is enabled on the network
- E. Unnecessary services are disabled on the server
- F. SELinux is enabled on the server



Suggested Answer: BC

Community vote distribution

BC (100%)

  **Ariel235788** Highly Voted 1 year, 9 months ago

IRL applications, I'd say more on B and C. Can't say adware for sure. Malware yes. but I would point more toward spyware if anything
upvoted 6 times

  **TheITStudent** Highly Voted 1 year, 4 months ago

Selected Answer: BC

To me, this seems like a clear case of C&C (command and control) someone hacked into the server, and it is now calling home. The hacker may have changed firewall rules to allow the server to call home, or they have been insecure to begin with. The problem is not the users computers having adware, the server is what is running slow, causing a denial of availability for the users. Server has malware. and server is calling home, a clear B & C for me. Have gotten Pentest+ Cysa+ Security+ Network+, most of the answers in these questions (by the community) are pretty bad. study this stuff for yourself, don't depend on the answers in the discussions. Just my opinion.
upvoted 5 times

  **Obi_Wan_Jacoby** Most Recent 11 months, 3 weeks ago



Selected Answer: BC

B&C are correct. "TheITStudent" explanation is right on.
upvoted 1 times

  **Rainkoot** 1 year, 4 months ago



Selected Answer: BC

"A server is performing slowly, ... several unauthorized services running on that server that are successfully communicating to an external site." End user issues are a by product of the server issues.
I agree with TheITStudent's explanation with C&C
upvoted 1 times

  **i_bird** 1 year, 4 months ago

Another Keyword CompTIA trick.. MOST like cause of the issue:

Adware is a type of malware, so if you choosing a virus as the possible cause you might as well choose adware as another possible cause, before misconfigured firewall.
upvoted 1 times

  **Dion79** 1 year, 7 months ago

why would firewall rules change? how is E the problem? harding is done before productions, so you have to re-harden your services... what if you shutdown the service to that application?

Why not check first to make sure workstations aren't infected with malware/adaware that could infect the server or spread through the network or to other users??..... Why is the server responding slowly and has external unknow communication..? Maybe workstations communicating with external port/application??

I like provided answers.

upvoted 1 times

🗨️ 👤 **PEsty93** 1 year, 8 months ago

Selected Answer: BC

I would say B and C

Adware on the users' machine is a large assumption. If everyone is affected it's unlikely that everyone has adware, and more likely that the server has a virus slowing it down.

The server most likely got the virus from open ports i.e. misconfigured firewall.

upvoted 1 times

🗨️ 👤 **dnc1981** 1 year, 9 months ago

C and E are correct

upvoted 2 times

🗨️ 👤 **[Removed]** 1 year, 8 months ago

C yes.

E. Unnecessary services are disabled on the server

Unnecessary - disabled

upvoted 1 times

A server technician is configuring the IP address on a newly installed server. The documented configuration specifies using an IP address of 10.20.10.15 and a default gateway of 10.20.10.254. Which of the following subnet masks would be appropriate for this setup?

- A. 255.255.255.0
- B. 255.255.255.128
- C. 255.255.255.240
- D. 255.255.255.254

Suggested Answer: A

Community vote distribution

A (100%)

- 🗳️ **slpcomputer** Highly Voted 3 years, 5 months ago

255.255.255.0

upvoted 7 times
- 🗳️ **Linenzo** Highly Voted 3 years, 5 months ago

10.20.10.15 and 10.20.10.254 should be in the same subnet

255.255.255.0 must be the subnet mask

upvoted 7 times
- 🗳️ **adadad5941** Most Recent 7 months, 2 weeks ago

255.255.255.0

upvoted 1 times
- 🗳️ **Obi_Wan_Jacoby** 2 years, 5 months ago

Selected Answer: A

Answer is A. If you do not know this, please take Net+ next.

upvoted 2 times
- 🗳️ **Pongsathorn** 2 years, 7 months ago

Selected Answer: A

255.255.255.0

upvoted 1 times
- 🗳️ **Rainkoot** 2 years, 10 months ago

Selected Answer: A

255.255.255.0 or 10.20.10.0/24

upvoted 2 times
- 🗳️ **Dion79** 3 years, 2 months ago

Selected Answer: A

I'd go with A.

upvoted 5 times
- 🗳️ **Dion79** 3 years, 1 month ago

specifies using an IP address of 10.20.10.15 and a default gateway of 10.20.10.254, would mean you would need 256 addresses, you would lose 2 addresses, giving you 254. So you would use CIDR /24 putting you at submask 255.255.255.0

upvoted 3 times
- 🗳️ **dnc1981** 3 years, 3 months ago


A is the correct answer

upvoted 2 times
- 🗳️ **PEsty93** 3 years, 3 months ago

Selected Answer: A

For both devices to be in the same subnet then it needs to include the entire range

upvoted 3 times

  **el_adamba** 3 years, 5 months ago

WTF? Do you really think that is right answer? I tried to use your answer and guess what?

Number of Usable Hosts: 0 - ROTFL

IPv4 Subnet Calculator

Result

IP Address: 10.20.10.15

Network Address: 10.20.10.14

Usable Host IP Range: NA



Broadcast Address: 10.20.10.15

Total Number of Hosts: 2

Number of Usable Hosts: 0

Subnet Mask: 255.255.255.254

upvoted 6 times

  **hisdayold112** 3 years, 2 months ago

you really think that is right answer? I tried to use your answer and guess what?

Number of Usable Hosts: 0 - ROTFL

IPv4 Subnet Calculator

Result

IP Address: 10.20.10.15

Network Address: 10.20.10.14

Usable Host IP Range: NA

Broadcast Address: 10.20.10.15

Total Number of Hosts: 2

Number of Usable Hosts: 0

Subnet Mask: 255.255.255.254

upvoted 1 times

A storage administrator is investigating an issue with a failed hard drive. A technician replaced the drive in the storage array; however, there is still an issue with the logical volume. Which of the following best describes the NEXT step that should be completed to restore the volume?

- A. Initialize the volume
- B. Format the volume
- C. Replace the volume
- D. Rebuild the volume

Suggested Answer: D

Community vote distribution

D (100%)

🗲️ 👤 **Obi_Wan_Jacoby** 11 months, 3 weeks ago

Selected Answer: D

I concur with D

upvoted 1 times

🗲️ 👤 **Dion79** 1 year, 8 months ago

<https://community.boschsecurity.com/t5/Security-Video/How-to-replace-defective-drive-rebuild-RAID-create-new-iSCSI/ta-p/45052>

upvoted 2 times

🗲️ 👤 **dnc1981** 1 year, 8 months ago

Selected Answer: D

Has to be D

upvoted 4 times

A large number of connections to port 80 is discovered while reviewing the log files on a server. The server is not functioning as a web server. Which of the following represent the BEST immediate actions to prevent unauthorized server access? (Choose two.)

- A. Audit all group privileges and permissions
- B. Run a checksum tool against all the files on the server
- C. Stop all unneeded services and block the ports on the firewall
- D. Initialize a port scan on the server to identify open ports
- E. Enable port forwarding on port 80
- F. Install a NIDS on the server to prevent network intrusions

Suggested Answer: CD

Community vote distribution

CD (100%)

🗳️ 👤 **Ariel235788** Highly Voted 2 years, 3 months ago

Selected Answer: CD

Auditing isn't preventative. It is Detective. CD are the correct answers.

upvoted 5 times

🗳️ 👤 **Ariel235788** 2 years, 3 months ago

You want to identify all open ports and disable the unneeded ones. Only way to do that is with a port scan. Running an audit is not an immediate preventative measure

upvoted 4 times

🗳️ 👤 **Pongsathorn** Most Recent 1 year, 7 months ago

Selected Answer: CD

Disable Unused Services/Close Unneeded Ports

Any services that are not required on the server should be disabled. Only those required for the server to perform its role in the network should be left on. The easiest way to do this is to install a host firewall on the system and adopt a "disable by default" policy with respect to services by closing the port used for the service. Then manually enable any you need.

upvoted 1 times

🗳️ 👤 **dddc1000** 1 year, 9 months ago

Okay, I think C and D. Here's why. For answer C, you can block port 80 at the firewall. Done! Now for answer D, running port scan on server will identify other unauthorized open ports. The answer can't be F because, it states NIDS, and the question is only focus on a single server. Which means if anything, you would install a HIDS to detect intrusions. Boom!

upvoted 1 times

🗳️ 👤 **Dion79** 2 years, 1 month ago

Lets talk about why the server is not functioning as a web server and what is it functioning as? File server? Directory Service? Doesn't state... Why is port 80 open if this server is not functioning as a web server? Maybe insider threat? rough employee? someone with elevated or admin rights messing with the server? horrible question and a trick.

upvoted 2 times

🗳️ 👤 **ITken** 9 months, 3 weeks ago

One of the tricks used on a compromised system is to utilize port 80 for communication. Regardless of what it's typically used for (http), any service can run on any port you configure it to run on. Threat actors will use it because it's a port that is commonly used for web traffic and, as such, will likely not be blocked by the firewall.

Now, on most computers, ports below 1024 are privileged ports that require an super user account for those services to bind to said ports. However, if the server is compromised, it's likely that the threat actors already have super user access to that system.

upvoted 1 times

🗳️ 👤 **Dion79** 2 years ago

Agree with others I'd go with C and D. A is definitely a possibility and CompTIA are the masters of word trickery.

upvoted 1 times

🗨️ 👤 **PEsty93** 2 years, 2 months ago

Selected Answer: CD

You would need to do A, to ensure nothing has been changed, but it won't prevent access and that is the question.

upvoted 2 times

🗨️ 👤 **dnc1981** 2 years, 3 months ago

C and maybe F are the immediate actions you would take. A is not an immediate action

upvoted 2 times

🗨️ 👤 **dnc1981** 2 years, 3 months ago

And D would take too long

upvoted 2 times

🗨️ 👤 **[Removed]** 2 years, 2 months ago

To my understanding not all networks/environments will have a NIDS. Port Sec is always a good idea and thing to check.


upvoted 1 times

A company is running an application on a file server. A security scan reports the application has a known vulnerability. Which of the following would be the company's BEST course of action?

- A. Upgrade the application package
- B. Tighten the rules on the firewall
- C. Install antivirus software
- D. Patch the server OS

Suggested Answer: A

Community vote distribution



🗳️ 👤 **Dion79** Highly Voted 2 years, 8 months ago

Selected Answer: A

first thing you would do.

upvoted 7 times

🗳️ 👤 **dnc1981** Highly Voted 2 years, 9 months ago

A is the obvious answer

upvoted 7 times

🗳️ 👤 **EngAbood** Most Recent 11 months, 4 weeks ago

Selected Answer: D

god help me with my choice /:

upvoted 1 times

🗳️ 👤 **Dingos** 1 year, 7 months ago

Selected Answer: D

I dont like word UPGRADE in answer A, what it means, does update even fix problem or make it even worse?

If they used patch or at least update.

I stick to D as "sometimes" server vendor found out some app weaknes-es add provide security for their OS. Microsoft for their OS also distribute patches for non MS software.

upvoted 1 times

🗳️ 👤 **Dopeboyroy** 1 year, 8 months ago

The answer is D. When you patch the server, you are updating the applications.

upvoted 1 times

🗳️ 👤 **kloug** 1 year, 9 months ago

aaaaaa

upvoted 1 times

🗳️ 👤 **Jmooney** 1 year, 9 months ago

Why would you UPGRADE the application package? If the word was "update" then the answer is A. But the only answer that has update is D...

upvoted 1 times

🗳️ 👤 **Pongsathorn** 2 years, 1 month ago

Selected Answer: A

Install Latest Patches

Don't forget about the applications that may be running on the server. Applications can also be attacked by hackers. That's why software vendors are also periodically issuing security updates. As security issues are reported, they respond by fixing the software. For Windows applications, these updates can accompany the operating system updates if you choose to enable them.

Other applications may be more of a challenge, but it's hard to find vendors today that don't either automatically send and install the updates or, at the very least, notify you that one is available.

upvoted 1 times

🗳️ 👤 **Dopeboyroy** 1 year, 8 months ago

How is the answer A., but the first thing you said was install the latest patches.

upvoted 1 times

🗨️ 👤 **[Removed]** 2 years, 8 months ago

Selected Answer: A

A. Happens regularly. Vulns discovered and install patch update.

upvoted 2 times

🗨️ 👤 **Ariel235788** 2 years, 8 months ago

Why would you replace the door if its your doorknob that is broken? replace the doorknob...

upvoted 6 times

🗨️ 👤 **i_bird** 2 years, 4 months ago

I meant have to replace my city, if my door knob is broken it's the BEST cause of action..@CompTIA

upvoted 1 times

A technician runs top on a dual-core server and notes the following conditions: top `"- 14:32:27, 364 days, 14 users load average 60.5 12.4 13.6
Which of the following actions should the administrator take?

- A. Schedule a mandatory reboot of the server
- B. Wait for the load average to come back down on its own
- C. Identify the runaway process or processes
- D. Request that users log off the server

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **gingasaurusrex** 8 months, 3 weeks ago

Selected Answer: C

the answer is C for sure

upvoted 1 times

🗳️ 👤 **kloug** 9 months, 1 week ago

no ccccccccccccccc correct

upvoted 1 times

🗳️ 👤 **kloug** 9 months, 1 week ago

aaaaaaaaaaaaaaa

upvoted 1 times

🗳️ 👤 **Laudy** 9 months, 4 weeks ago

Selected Answer: C

Well, you don't want to reboot unless it's absolutely necessary. So no. And why would I boot users who have nothing to do with the problem? Find the problem process, and maybe the associated user, and address it from there.

upvoted 2 times

🗳️ 👤 **Pongsathorn** 1 year, 1 month ago

Selected Answer: C

Load Average in Linux is a metric that is used by Linux users to keep track of system resources. It also helps you monitor how the system resources are engaged.

To understand the Load Average in Linux, we need to know what do we define as load. In a Linux system, the load is a measure of CPU utilization at any given moment.

It refers to the number of processes which are either currently being executed by the CPU or are waiting for execution.

An idle system has a load of 0. With each process that is being executed or is on the waitlist, the load increases by 1.

Occasionally a process will stop responding to the system and run wild. These processes ignore their scheduling priority and insist on taking up 100% of the CPU. Because other processes can only get limited access to the CPU, the machine begins to run very slowly.

<https://www.digitalocean.com/community/tutorials/load-average-in-linux>

<https://docs.cs.byu.edu/doku.php?id=runaway-processes>

upvoted 2 times

🗳️ 👤 **Pongsathorn** 1 year, 1 month ago

(cont.)



How can I identify a runaway process on my computer?

The 'w' command at the terminal will print out a list of current users of a machine, and it will tell you the machine's "load average." The load average of a machine is related to how much input/output the machine has to do. A load average of 1 is a machine under full load. Anything over 1 is extremely high and means that the machine is getting behind on its processing. If your machine has a load average near or over 1, and you are not running anything really resource intensive on the machine, then you probably have a runaway process sapping your machine's processing power.

<https://www.digitalocean.com/community/tutorials/load-average-in-linux>

<https://docs.cs.byu.edu/doku.php?id=runaway-processes>

upvoted 2 times

  **Timock** 1 year, 3 months ago

Users logging off does not seem to be the answer here. I would identify the processes and possibly reboot the server depending on what I find going wrong with the TOP command.

upvoted 2 times

A technician needs to set up a server backup method for some systems. The company's management team wants to have quick restores but minimize the amount of backup media required. Which of the following are the BEST backup methods to use to support the management's priorities? (Choose two.)

- A. Differential
- B. Synthetic full
- C. Archive
- D. Full
- E. Incremental
- F. Open file

Suggested Answer: AB

Community vote distribution

BE (67%)

AB (33%)

🗳️ 👤 **Melendez** 2 months, 1 week ago

Selected Answer: AB

Chatgpt

upvoted 1 times

🗳️ 👤 **broman** 9 months, 3 weeks ago

The answers are AB check out this link <https://aws.amazon.com/compare/the-difference-between-incremental-differential-and-other-backups/#:~:text=A%20differential%20backup%20strategy%20only,changes%20since%20the%20last%20backup.>

upvoted 1 times

🗳️ 👤 **a792193** 1 year ago

Selected Answer: AB

An incremental backup is a backup of changes since a specified marker (usually a full backup) and resets the archive bit. Backups are quicker than differentials, but the restoral process is slower (must use each backup job after the full).

This means we want differential and synthetic full

upvoted 2 times

🗳️ 👤 **Mareo** 1 year, 10 months ago

Selected Answer: BE

Incremental: This method backs up only the changes since the last backup, which minimizes the amount of backup media required and can be performed quickly.

Synthetic full: This method creates a full backup by combining the last full backup with the changes since the last backup, which minimizes the amount of backup media required and provides a quick restore.

upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 1 year, 11 months ago

Selected Answer: BE

I am going with B&E

upvoted 1 times

🗳️ 👤 **Pongsathorn** 2 years, 1 month ago

Differential: Takes less space than full backups, Faster restoration than incremental backups

Synthetic full

Full: Quick restore time, take up large amounts of unnecessary storage space

Incremental: Time-consuming restoration since data must be pieced together from multiple backups

<https://www.unitrends.com/blog/types-of-backup-full-incremental-differential>

<https://gomindsight.com/insights/blog/four-data-backup-methods-it-resiliency/>

upvoted 2 times

🗨️ 👤 **Don_Nguy3n** 2 years, 5 months ago

Key word 'but minimize the amount of backup media required'

So It Incremental and synthetic

upvoted 2 times

🗨️ 👤 **Dion79** 2 years, 8 months ago

I'd go with provided answers.

Full Slow backup, fast restore

Full+Incremental Fast backup, slow restore (compared to Full+Differential)

Full+Differential Slow backup, fast restore (compared to Full+Incremental)

Reference: The Official CompTIA Server+ Study Guide (Exam SK0-005)

upvoted 3 times

🗨️ 👤 **dnc1981** 2 years, 9 months ago

B is quickest and E would take the least amount of time

upvoted 2 times

🗨️ 👤 **Dion79** 2 years, 8 months ago

E - Incremental would increase the amount of media and take longer time to restore. Faster backup time.

A - Differential would decrease the amount of media and take shorter time to restore. Longer backup time.

B - Synthetic Full - fast restore

Synthetic Full/Differential meets managements requirements.

upvoted 2 times

🗨️ 👤 **dnc1981** 2 years, 8 months ago

Differential (and Incremental) would actually *increase* the amount of tapes needed in the event of a restore, because you'd have to restore the most recent Full backup as well as the most recent Differential (or in the case of Incrementals, all incrementals since the most recent Full)

upvoted 1 times

Ann, an administrator, is configuring a two-node cluster that will be deployed. To check the cluster's functionality, she shuts down the active node. Cluster behavior is as expected, and the passive node is now active. Ann powers on the server again and wants to return to the original configuration. Which of the following cluster features will allow Ann to complete this task?

- A. Heartbeat
- B. Failback
- C. Redundancy
- D. Load balancing

Suggested Answer: B

  **Dion79**  10 months, 1 week ago

When you research server clustering, you frequently hear the terms failover and failback. Failover refers to the failure of the active service provider node and passive node's taking over of its responsibilities. At that point, the server that was the passive node becomes the active node. It services clients and handles the cluster's duties. Once the original active node returns to service, you may set it to act as the new passive node, or you may shift the services back to it, causing it to take back its active node role and relegating the formerly passive server to its original role. The term failback refers to the service returning to the original active node.

upvoted 21 times

Which of the following policies would be BEST to deter a brute-force login attack?

- A. Password complexity
- B. Password reuse
- C. Account age threshold
- D. Account lockout threshold

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **ompk** 9 months ago

I wonder who is choosing the wrong answers in this website for this course?
upvoted 1 times

🗳️ 👤 **tame_rabbit** 1 year ago

Selected Answer: D

The best policy to deter a brute-force login attack is D. Account lockout threshold. This policy directly limits the number of login attempts, thereby preventing an attacker from continuously attempting to guess the correct password. Implementing an account lockout threshold makes brute-force attacks impractical by significantly increasing the time required to successfully guess a password.
upvoted 2 times

🗳️ 👤 **Mareo** 2 years, 5 months ago

Selected Answer: D

Because a policy of locking out accounts after a certain number of failed login attempts can effectively deter brute-force attacks. This is because the attacker will be unable to continue trying different passwords once the account has been locked out, making it more difficult to gain unauthorized access.
upvoted 1 times

🗳️ 👤 **nixonbii** 2 years, 8 months ago

Selected Answer: D

Attacker cannot continue to try different passwords if the account gets locked out. That's actually from A+.
upvoted 2 times

🗳️ 👤 **Fineb** 2 years, 9 months ago

Dis the correct answer, locking out the user or attacker after few tries will prevent that
upvoted 1 times

🗳️ 👤 **jagoichi** 2 years, 9 months ago

D Lockout policy is the BEST option
The other options create a strong password policy
upvoted 2 times

A technician needs to install a Type 1 hypervisor on a server. The server has SD card slots, a SAS controller, and a SATA controller, and it is attached to a NAS.

On which of the following drive types should the technician install the hypervisor?

- A. SD card
- B. NAS drive
- C. SATA drive
- D. SAS drive

Suggested Answer: D

Community vote distribution



dnc1981 3 years, 3 months ago

D is the correct answer. SATA would not be used on a server because its not server grade storage
upvoted 8 times

TheITStudent 2 years, 10 months ago

Selected Answer: A

I agree with Dion79... going with SD card. Specifically because the questions states that the server has sd card slot available, in that case, these are there FOR THAT PARTICULAR PURPOSE... see: <https://serverfault.com/questions/823791/what-is-the-use-case-of-a-sd-slot-on-motherboard>
upvoted 5 times

Kabbath1986 1 month, 4 weeks ago

Selected Answer: A

✓ A. SD card

Explanation:

A Type 1 hypervisor (bare-metal hypervisor) like VMware ESXi or Microsoft Hyper-V is installed directly on the hardware. It typically requires very little disk space and minimal IOPS since the actual virtual machines and their workloads are run from separate storage.

SD cards are commonly used by OEMs (like Dell, HPE, etc.) to host hypervisors because:

The hypervisor has low write requirements.

It leaves SATA and SAS drives free for storing virtual machines.

They are cost-effective for this use.

Note: If persistent logging or frequent writes are required, some admins configure logs to go to other storage (e.g., a data store on SAS or SATA) to avoid wearing out the SD card.

upvoted 1 times

PaytoPlay 3 months, 1 week ago

Selected Answer: D

Going with SAS over SATA here, I mean you have it so why not?

upvoted 1 times

DudeGuy13 4 months ago

Selected Answer: D

SD cards typically do not support RAID configurations. A tenant of server management is redundancy.

upvoted 1 times

cj207800 5 months ago

Selected Answer: A

I agree with Dion79... going with SD card. An SD card is an ideal choice for installing a Type 1 hypervisor because. It provides a dedicated, local boot device for the hypervisor.

It separates the hypervisor from the main storage, allowing for easier management and upgrades.

Many enterprise-grade servers come with internal SD card slots specifically designed for hypervisor installation.

While SATA or SAS drives could technically be used, they are better suited for VM storage and data. The NAS drive is not a suitable option for a Type 1 hypervisor installation, as it's not directly attached to the server hardware and could introduce performance and reliability issues

upvoted 1 times

🗳️ 👤 **ibn_e_nazir** 5 months, 1 week ago

Selected Answer: D

Option D SAS DRIVE is the correct answer. The question clearly says that "On which of the following drive types should the technician install the hypervisor?". SD card is not a drive type and it is not normally used as source to run an OS. Running a Type 1 hypervisor on an SD card in a production environment is not ideal. They can be used to host a lightweight embedded system.

upvoted 1 times

🗳️ 👤 **kx7tg4xu** 10 months, 1 week ago

Selected Answer: C

C is the correct answer.

In many common environments, a hypervisor can be run on SATA without any problems.

Hypervisors themselves usually do not require large storage capacities or extreme speeds, so SATA drives are often sufficient.

upvoted 1 times

🗳️ 👤 **ccoli** 1 year, 1 month ago

It is D, if they said it had SD card slots with SD cards installed it would be A but it has a SAS controller attached to a NAS and empty SD cards so you use SAS.

upvoted 1 times

🗳️ 👤 **AzadOB** 1 year, 3 months ago

Selected Answer: A

For installing a Type 1 hypervisor on a server, the technician should install it on:

A. SD card

Explanation:

Installing the hypervisor on an SD card is a common practice for Type 1 hypervisors. SD cards offer a convenient and cost-effective solution for booting the hypervisor directly from the server hardware. This approach allows for efficient utilization of internal storage resources while keeping the hypervisor separate from the storage used for virtual machines.

upvoted 1 times

🗳️ 👤 **kloug** 2 years, 3 months ago

aaaaaaa

upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 2 years, 5 months ago

Selected Answer: D

D (as others mentioned) is correct

upvoted 2 times

🗳️ 👤 **momoci** 2 years, 7 months ago

Selected Answer: D

Agree with Dion79 Correct Answer is D

upvoted 1 times

🗳️ 👤 **paperburn** 2 years, 8 months ago

Selected Answer: D

Assuming SD Slots are internal. If so, then Internal SD Slots are recommended.

upvoted 2 times

🗳️ 👤 **szl0144** 2 years, 9 months ago

Selected Answer: D

D is the answer

upvoted 1 times

🗳️ 👤 **Dion79** 3 years, 1 month ago

A & D are both possible answers. All components described in this example are internal. Assuming SD Slots are internal. If so, then Internal SD Slots are recommended. "The Internal SanDisk (SD) slots remain enabled because the ESXI hypervisor runs on dual SD cards inside the server chassis".

Reference: Lesson 9 under Configure Server Hardening. 31 of 93 CompTIA Server+ Final Assessment.

upvoted 2 times

A technician is trying to determine the reason why a Linux server is not communicating on a network. The returned network configuration is as follows: `eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500 inet 127.0.0.1 network 255.255.0.0 broadcast 127.0.0.1`
Which of the following BEST describes what is happening?

- A. The server is configured to use DHCP on a network that has multiple scope options
- B. The server is configured to use DHCP, but the DHCP server is sending an incorrect subnet mask
- C. The server is configured to use DHCP on a network that does not have a DHCP server
- D. The server is configured to use DHCP, but the DHCP server is sending an incorrect MTU setting

Suggested Answer: C

Community vote distribution


C (75%)

B (25%)

  **dnc1981**  3 years, 3 months ago

I think it's C

upvoted 7 times

  **Pongsathorn** 2 years, 7 months ago

If the network does not have a DHCP server then the server should get APIPA that self-assigned address when the client fails to lease an address from a DHCP server.

upvoted 1 times

  **[Removed]** 10 months, 2 weeks ago

Not for Linux

upvoted 2 times

  **[Removed]**  10 months, 3 weeks ago

Selected Answer: C

It defaulted to a loopback, and it's a Linux box. There is just 0 reason why someone would setup a DHCP to assign a loopback address, even with the unusual subnet mask.

Even ChatGPT seems to agree, after I asked about the subnet mask, the loopback instead of apiipa (which by the way is avahi on linux) and despite all that, seemed to think that C was the better answer.

C makes more sense than someone having setup a DHCP server with that address and getting an incorrect subnet mask.

upvoted 2 times

  **tame_rabbit** 1 year, 2 months ago

Selected Answer: C

It's evident that the server's network interface eth0 is configured with the IP address 127.0.0.1, which is the loopback address. This address is reserved for communication within the local system and is not used for external network communication. Additionally, the subnet mask 255.255.0.0 is associated with this loopback address.

Given this information, the most suitable explanation is:

C. The server is configured to use DHCP on a network that does not have a DHCP server.

The Server will not get an APIPA address because this is a Linux server

upvoted 1 times

  **Pongsathorn** 2 years, 7 months ago

Selected Answer: B



If the network does not have a DHCP server then the server should get APIPA that self-assigned address when the client fails to lease an address from a DHCP server.

upvoted 1 times

  **[Removed]** 10 months, 2 weeks ago

Not for Linux

upvoted 1 times

  **King2** 2 years, 9 months ago

I think provided answer (B) is correct

The loopback interface should be assigned a netmask of 255.0.0.0, since 127.0.0.1 is a class A address.

Ref: <https://www.oreilly.com/library/view/linux-network-administrators/1565924002/ch05s07.html>

upvoted 2 times

A server technician is deploying a server with eight hard drives. The server specifications call for a RAID configuration that can handle up to two drive failures but also allow for the least amount of drive space lost to RAID overhead. Which of the following RAID levels should the technician configure for this drive array?

- A. RAID 0
- B. RAID 5
- C. RAID 6
- D. RAID 10

Suggested Answer: C

Reference:

<https://www.booleanworld.com/raid-levels-explained/>

Community vote distribution

C (100%)

🗨️ 👤 **Pongsathorn** 7 months, 2 weeks ago

Selected Answer: C

The RAID 6 design distributes data across a minimum of four HDDs the same way that RAID 0 and RAID 5 do, but it also distributes parity information across two disks. The result is that a RAID 6 array can recover data even with the failure of two HDDs. Reads are quick, like RAID 5, but writes are slower due to the duplication of the parity data. It is a good general solution as long as the performance hit on write tasks is not a problem for your environment.

Ref. The Official CompTIA Server+ Study Guide (Exam SK0-005) page 114.

upvoted 3 times

🗨️ 👤 **Dion79** 1 year, 4 months ago

RAID 6 Disk Striping with Double Parity

The RAID 6 design distributes data across a minimum of four HDDs the same way that RAID 0 and RAID 5 do, but it also distributes parity information across two disks. The result is that a RAID 6 array can recover data even with the failure of two HDDs. Reads are quick, like RAID 5, but writes are slower due to the duplication of the parity data. It is a good general solution as long as the performance hit on write tasks is not a problem for your environment.

RAID 6 requires a minimum of four HDDs.

upvoted 4 times

Which of the following should an administrator use to transfer log files from a Linux server to a Windows workstation?

- A. Telnet
- B. Robocopy
- C. XCOPY
- D. SCP

Suggested Answer: D

Community vote distribution

D (100%)

🗲️ 👤 **Dion79** Highly Voted 1 year, 3 months ago

The WinSCP tool is commonly used to transfer files between Linux and Windows systems. WinSCP has an easy GUI and encrypts the contents of the file transfer.

upvoted 8 times

🗲️ 👤 **Pongsathorn** Most Recent 7 months, 2 weeks ago

Selected Answer: D

SCP

Linux also relies on scp to provide encrypted file transfers using SSH. The Linux command to copy a file is cp. Therefore, scp is secure copy. It uses the same encryption mechanism as SSH to guarantee file integrity and confidentiality.

If there is WinSCP available, WinSCP is the best answer.

The WinSCP tool is commonly used to transfer files between Linux and Windows systems. WinSCP has an easy GUI and encrypts the contents of the file transfer.

Ref. The Official CompTIA Server+ Study Guide (Exam SK0-005) page 198.

upvoted 3 times

🗲️ 👤 **Pongsathorn** 7 months, 3 weeks ago

Selected Answer: D

Xcopy and Robocopy are two Windows built-in command line file copy utilities. Both of them can help you copy files and folders from one location to another.

upvoted 1 times

🗲️ 👤 **TheITStudent** 10 months, 2 weeks ago

Selected Answer: D

https://winscp.net/eng/docs/transfer_mode

upvoted 1 times

🗲️ 👤 **TheITStudent** 10 months, 2 weeks ago



Also, robocopy & xcopy are windows os only commands for windows machines, telnet is irrelevant

upvoted 2 times

Users in an office lost access to a file server following a short power outage. The server administrator noticed the server was powered off. Which of the following should the administrator do to prevent this situation in the future?

- A. Connect the server to a KVM
- B. Use cable management
- C. Connect the server to a redundant network
- D. Connect the server to a UPS

Suggested Answer: *D*

  **Kraken84** 8 months, 2 weeks ago

DDDD.. doi

upvoted 1 times

Which of the following describes the installation of an OS contained entirely within another OS installation?

- A. Host
- B. Bridge
- C. Hypervisor
- D. Guest

Suggested Answer: D

Community vote distribution

D (100%)

  **Dion79** Highly Voted 1 year, 6 months ago

Selected Answer: D

Guest is the answer of the riddle.

upvoted 7 times

  **Ariel235788** Highly Voted 1 year, 8 months ago

Why wouldn't it be Guest?

upvoted 5 times

  **dnc1981** 1 year, 8 months ago

It's definitely guest. The answers on this site are sometimes wrong

upvoted 3 times

  **Dingos** Most Recent 7 months, 3 weeks ago

Selected Answer: D

I am confused how it can be Hypervisor

upvoted 1 times

  **momoci** 1 year, 1 month ago

Selected Answer: D

The key point on this Question is OS installation that contained entirely within other OS. The Answer is Guest

upvoted 1 times

  **Pongsathorn** 1 year, 1 month ago

Selected Answer: D

Host vs. Guest

The foundation of virtualization is the host device, which may be a workstation or a server. This device is the physical machine that contains the software that makes virtualization possible and the containers, or virtual machines, for the guest operating systems. The host provides the underlying hardware and computing resources, such as processing power, memory, disk, and network I/O, to the VMs. Each guest is a completely separate and independent instance of an operating system and application software.

The host is responsible for allocating compute resources to each of the VMs as specified by the configuration. The software that manages all of this is called the hypervisor. Based on parameters set by the administrator, the hypervisor may take various actions to maintain the performance of each guest as specified by the administrator

upvoted 4 times

  **paperburn** 1 year, 2 months ago

Selected Answer: D

Guest is the answer

upvoted 2 times

  **cbrp** 1 year, 3 months ago

Selected Answer: D

Guest is the answer

upvoted 3 times

A server technician is installing a Windows server OS on a physical server. The specifications for the installation call for a 4TB data volume. To ensure the partition is available to the OS, the technician must verify the:

- A. hardware is UEFI compliant
- B. volume is formatted as GPT
- C. volume is formatted as MBR
- D. volume is spanned across multiple physical disk drives

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **kloug** 9 months, 1 week ago

bbbbbbbbb

upvoted 1 times

🗳️ 👤 **momoci** 1 year, 1 month ago

Selected Answer: B

Since when partition volume is about UEFI?? Of course The Answer is B vote for Pongsathorn.

upvoted 2 times

🗳️ 👤 **Pongsathorn** 1 year, 1 month ago

Selected Answer: B

PARTITION TABLES

Master Boot Record (MBR) and GUID Partition Table (GPT)

HDDs are partitioned to organize data, and the location of these partitions must be maintained. There are two different types of tables used to relate partition locations on storage disks. The older method is the MBR, and the newer way is the GPT. The primary difference between the two is that the GPT is far more flexible and practical on modern servers.+

GPT

- Supports a larger number of partitions on the HDD.
- Recognizes drives that are larger than two TB.

MBR

- Supports only four partitions on the HDD.
- Recognizes drives that are two TB or smaller.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 134.

upvoted 3 times

🗳️ 👤 **Pongsathorn** 1 year, 1 month ago

(cont.)

Modern servers will likely exceed both of these requirements, and therefore a GPT configuration is the best bet. Server firmware must support UEFI system configurations to utilize a GPT structure.

Older servers may be configured with the MBR if their drive space is more limited or they need fewer partitions. MBR cannot manage more than 2 TB of storage space.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 134.

upvoted 1 times

🗳️ 👤 **Pongsathorn** 1 year, 2 months ago

Selected Answer: B

MBR works with disks up to 2 TB in size, but it can't handle larger disks. MBR also supports only up to four primary partitions, so to have more than four, you had to make one of your primary partitions an "extended partition" and create logical partitions inside it. GPT removes both of these limitations. It allows up to 128 partitions on a GPT drive.

upvoted 1 times

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

GPT is also used on some BIOS systems because of the limitations of MBR partition tables, which was the original driver for the development of UEFI/GPT. MBR works with disks up to 2 TB in size, but it can't handle larger disks. MBR also supports only up to four primary partitions, so to have more than four, you had to make one of your primary partitions an "extended partition" and create logical partitions inside it. GPT removes both of these limitations. It allows up to 128 partitions on a GPT drive.

upvoted 1 times

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

UEFI is a standard firmware interface for servers and PCs designed to replace BIOS. Here are some advantages of UEFI firmware:

- Better security; protects the preboot process
- Faster startup times and resuming from hibernation
- Support for drives larger than 2.2 terabytes (TB)
- Support for 64-bit firmware device drivers
- Capability to use BIOS with UEFI hardware

upvoted 1 times

🗨️ 👤 **jagoichi** 1 year, 3 months ago

GPT is the answer

reference Official comptia server +

GPT Supports a larger number of partitions on the HDD

Recognizes drives that are larger than two TB

MBR Supports only four partitions on the HDD

Recognizes drives that are two TB or smaller

Older servers may be configured with the MBR if their drive space is more limited or they need fewer partitions. MBR cannot manage more than 2 TB of storage space

upvoted 1 times

🗨️ 👤 **PEsty93** 1 year, 8 months ago

Selected Answer: B

Answer is B. While UEFI is required to boot from a volume more than 2TB. GPT would be required for a data volume.

upvoted 2 times

🗨️ 👤 **i_bird** 1 year, 4 months ago

GPT was developed to replace the limitations of MBR, and it is part of UEFI development to replace PC BIOS.

So the Hardware has to UEFI compatible for GPT to be fully effective.

upvoted 2 times

🗨️ 👤 **dnc1981** 1 year, 9 months ago

B is the correct answer

upvoted 2 times

An administrator is configuring a server that will host a high-performance financial application. Which of the following disk types will serve this purpose?

- A. SAS SSD
- B. SATA SSD
- C. SAS drive with 10000rpm
- D. SATA drive with 15000rpm

Suggested Answer: A

Reference:

<https://www.hp.com/us-en/shop/tech-takes/sas-vs-sata>

1. The bytes that make up your document are saved in storage.
2. The computer signals to storage that it needs the bytes.
3. The bytes leave storage via the SAS and SATA cables and travel to the CPU at the motherboard - it's the CPU that runs Microsoft Word.
4. When you've finished writing for the day, you save the file.
5. The bytes travel back along the SAS or SATA cables and reenter storage.

Community vote distribution

A (100%)

🗳️ 👤 **Katlegobogosi** 10 months, 1 week ago

Selected Answer: A

SAS SSD

upvoted 1 times

🗳️ 👤 **Pongsathorn** 1 year, 1 month ago

Selected Answer: A

Comparison of the three DAS implementations:

- SATA when cost is more important than performance, less need for scalability
- SAS to balance cost and performance, scalability
- NVMe when performance is more important than cost, especially on servers that see large file transfers

Ref. The Official CompTIA Server+ Study Guide (Exam SK0-005) page 112.

upvoted 1 times

Which of the following DR testing scenarios is described as verbally walking through each step of the DR plan in the context of a meeting?

- A. Live failover
- B. Simulated failover
- C. Asynchronous
- D. Tabletop

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ **haazybanj** Highly Voted 2 years, 2 months ago

Tabletop is more appropriate
upvoted 6 times

🗳️ **Alfred69** Most Recent 9 months, 4 weeks ago

D. TableTop/Simulated Failover - the DRP are implemented on a limited scale. Participants engage in role-playing to ensure comprehension and realism.
upvoted 1 times

🗳️ **gingasaurusrex** 1 year, 2 months ago

Selected Answer: D
it is def table top
upvoted 1 times

🗳️ **kloug** 1 year, 3 months ago

dddddddddd
upvoted 1 times

🗳️ **Katlegobogosi** 1 year, 4 months ago

Selected Answer: D
D. Tabletop is the DR testing scenario that is described as verbally walking through each step of the DR plan in the context of a meeting. This type of test is often used to evaluate the effectiveness of a disaster recovery plan and to identify areas for improvement. It involves a group of stakeholders discussing hypothetical scenarios and the steps they would take to respond to them, without actually executing any of the recovery procedures.
upvoted 1 times

🗳️ **Pongsathorn** 1 year, 8 months ago

Selected Answer: D
Tabletops
Conducting a tabletop exercise is the most cost-effective and efficient way to identify areas of vulnerability before moving on to more involved testing. A tabletop exercise is an informal brainstorming session that encourages participation from business leaders and other key employees. In a tabletop exercise, the participants agree to determine a particular attack scenario upon which they then focus.
CompTIA Server+ Study Guide: Exam SK0-005 Chapter 9 Disaster Recovery
upvoted 1 times

🗳️ **nixonbii** 1 year, 8 months ago

Selected Answer: D
Answer is D. A live failover is a real scenario in which an active/active cluster configuration experiences a failure in one of the nodes which results in zero downtime due to the active redundancy of the other nodes.
upvoted 1 times

🗳️ **paperburn** 1 year, 8 months ago

Selected Answer: D
Tabletop/Simulated failover—the disaster recovery procedures are implemented on a limited scale. Participants engage in role-playing to ensure comprehension and realism.
Parallel recovery using a non-production test environment—the disaster recovery procedures are implemented in a non-production environment. VMs work especially well for this kind of test. The production environment remains unaffected.



upvoted 1 times

  **paperburn** 1 year, 8 months ago

Tabletop/Simulated failover—the disaster recovery procedures are implemented on a limited scale. Participants engage in role-playing to ensure comprehension and realism.



Parallel recovery using a non-production test environment—the disaster recovery procedures are implemented in a non-production environment. VMs work especially well for this kind of test. The production environment remains unaffected.

upvoted 1 times

  **Dion79** 2 years, 4 months ago

Thinks this one is wrong... Sounds more like Tabletop or even walk through

upvoted 4 times

  **Dion79** 2 years, 4 months ago

Paper test—critical stakeholders examine the disaster recovery procedures in the organization, and suggestions are considered.

Walk-through—the disaster recovery procedures are stepped through to confirm their viability. No changes are made, and no data is restored.

Tabletop/Simulated failover—the disaster recovery procedures are implemented on a limited scale. Participants engage in role-playing to ensure comprehension and realism.

Parallel recovery using a non-production test environment—the disaster recovery procedures are implemented in a non-production environment. VMs work especially well for this kind of test. The production environment remains unaffected.

Live failover (cutover)—the disaster recovery procedure is tested on the production environment where customers and employees reside. The goal is to prove zero interruption of service.

upvoted 3 times

When configuring networking on a VM, which of the following methods would allow multiple VMs to share the same host IP address?

- A. Bridged
- B. NAT
- C. Host only
- D. vSwitch

Suggested Answer: B

Reference:

<https://www.virtualbox.org/manual/ch06.html>

Table 6.1. Overview of Networking Modes

Mode	VM→Host	VM←Host	VM1↔VM2	VM→Net/LAN	VM←Net/LAN
Host-only	+	+	+	-	-
Internal	-	-	+	-	-
Bridged	+	+	+	+	+
NAT	+	Port forward	-	+	Port forward
NATservice	+	Port forward	+	+	Port forward

Community vote distribution

B (100%)

 **Katlegobogosi** 10 months, 1 week ago

Selected Answer: B

network address translation

upvoted 1 times

 **Obi_Wan_Jacoby** 11 months, 3 weeks ago

I concur with B

upvoted 1 times

 **paperburn** 1 year, 2 months ago

Nat Network attached storage

upvoted 1 times

 **2323323232** 1 year ago

network address translation (NAT) A way of remapping one IP address space into another. The most common use is to enable internal hosts to gain Internet access where the source IP address for outgoing transmissions is translated to the NAT router's public interface IP address.

upvoted 1 times

A technician recently upgraded several pieces of firmware on a server. Ever since the technician rebooted the server, it no longer communicates with the network.

Which of the following should the technician do FIRST to return the server to service as soon as possible?

- A. Replace the NIC
- B. Make sure the NIC is on the HCL
- C. Reseat the NIC
- D. Downgrade the NIC firmware

Suggested Answer: D

Community vote distribution



🗳️ 👤 **kx7tg4xu** 10 months, 1 week ago

Selected Answer: C

I think C is correct.

"C. Reseat the NIC " is the best solution to try first. This is the easiest and quickest method, as it will most likely solve the physical connection problem. The other options are next-stage measures that should be considered after trying this simple procedure.

This approach follows the basic principle of problem solving: start with the simplest and quickest solution. It is also most suited to the "recover the server as quickly as possible" requirement.

upvoted 2 times

🗳️ 👤 **RBL23168** 1 year, 7 months ago

Sounds to me like they updated drivers or something like that that didn't install properly. D first, then A if that doesn't resolve it.

upvoted 2 times

🗳️ 👤 **peachycaper** 1 year, 8 months ago

"several pieces of firmware" Firmware is not a physical device, no new hardware was installed so answer is D, basically roll it back.

upvoted 2 times

🗳️ 👤 **Dingos** 2 years, 1 month ago

Selected Answer: B

Why? Because:

"Checking the HCL is an important step when upgrading firmware, installing new hardware components, or updating software versions. It helps to ensure that the system is running on compatible hardware and software, which can minimize compatibility issues, performance problems, and system instability."

upvoted 1 times

🗳️ 👤 **Riseofashes** 2 years, 2 months ago

Selected Answer: D

If this was the result of a NIC replacement, reseating it would be the first choice.

However this appears to be a firmware issue, and downgrading it again would be the quickest way to bring it back online.

upvoted 4 times

🗳️ 👤 **kloug** 2 years, 3 months ago

ccccccccc

upvoted 1 times

🗳️ 👤 **Katlegobogosi** 2 years, 4 months ago

Selected Answer: C

Before downgrading the NIC firmware, the technician should try reseating the NIC to ensure it is properly connected to the motherboard. This is a simple and quick solution that may resolve the issue. If reseating the NIC does not resolve the problem, the technician may need to consider other troubleshooting options, such as checking the NIC compatibility with the server or replacing the NIC if necessary.

upvoted 2 times

🗳️ 👤 **lborden** 2 years, 4 months ago



It states he upgraded several pieces of firmware, so why wouldn't you downgrade the NIC firmware first? He didn't install a new NIC or touch the hardware, right?

upvoted 3 times

  **Replicant** 2 years, 3 months ago

All he did was reboot and now it doesn't work. Reseating the card is not a logical first step.

upvoted 2 times

  **nixonbii** 2 years, 8 months ago

Too ambiguous to be able to answer. Too many assumptions need to be made by the test taker.

upvoted 1 times

A server administrator has noticed that the storage utilization on a file server is growing faster than planned. The administrator wants to ensure that, in the future, there is a more direct relationship between the number of users using the server and the amount of space that might be used. Which of the following would BEST enable this correlation?

- A. Partitioning
- B. Deduplication
- C. Disk quotas
- D. Compression

Suggested Answer: C

Community vote distribution

C (100%)

🗲️ 👤 **Katlegobogosi** 10 months, 1 week ago

Selected Answer: C

C. Disk quotas

upvoted 1 times

🗲️ 👤 **Obi_Wan_Jacoby** 11 months, 3 weeks ago

Selected Answer: C

Obviously C

upvoted 1 times

A server administrator needs to keep a copy of an important fileshare that can be used to restore the share as quickly as possible. Which of the following is the BEST solution?

- A. Copy the fileshare to an LTO-4 tape drive
- B. Configure a new incremental backup job for the fileshare
- C. Create an additional partition and move a copy of the fileshare
- D. Create a snapshot of the fileshare

Correct Answer: *D*

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

Which of the following can be BEST described as the amount of time a company can afford to be down during recovery from an outage?

- A. SLA
- B. MTBF
- C. RTO
- D. MTTR

Suggested Answer: C

Community vote distribution

C (100%)

 **Pongsathorn** Highly Voted 7 months, 3 weeks ago

Selected Answer: C

Recovery Time Objective (RTO)

This is the shortest time period after a disaster or disruptive event within which a resource or function must be restored in order to avoid unacceptable consequences. RTO assumes that an acceptable period of downtime exists.

upvoted 6 times

Which of the following actions should a server administrator take once a new backup scheme has been configured?

- A. Overwrite the backups
- B. Clone the configuration
- C. Run a restore test
- D. Check the media integrity

Suggested Answer: D

Community vote distribution

C (50%)

D (50%)

PEsty93 Highly Voted 3 years, 2 months ago

I think D is correct. Only a scheme has been created. No backups have been taken yet so you cannot so a test restore.
upvoted 11 times

haazybanj Highly Voted 3 years, 2 months ago

This should be C
upvoted 7 times

Dion79 3 years, 1 month ago

I'd go with C as well. Running a restore will validate integrity and if backup was successful.
upvoted 1 times

Kabbath1986 Most Recent 1 month, 1 week ago

Selected Answer: C

The correct answer is: C. Run a restore test

Explanation:

Once a new backup scheme has been configured, the most critical step is to verify that you can successfully restore data from the backup. Backups are only useful if they can be restored reliably when needed.

✓ Why C. Run a restore test is correct:

It validates the backup process.

Ensures that data can be recovered in case of an actual failure or disaster.

Helps identify misconfigurations in the backup process early.

upvoted 1 times

Fart2023 11 months, 1 week ago

Selected Answer: C

By doing C you are testing D and more.
upvoted 1 times

RBL23168 1 year, 6 months ago

Selected Answer: D

. If a new backup scheme has been configured but no actual backups have been taken yet, then checking the media integrity (option D) would indeed be a suitable action to take. This ensures that the backup media (such as tapes, disks, etc.) is reliable and can be trusted for future backups. Running a restore test (option C) would typically be done after actual backups have been taken to ensure that the restoration process works correctly.

upvoted 3 times

MrS 1 year, 12 months ago

Selected Answer: C

Once a new backup scheme has been configured, a server administrator should run a restore test. A restore test involves restoring data from the backup to ensure that the backup is working correctly and that the data can be successfully restored in the event of a disaster. Running a restore test is an important step in verifying that the backup scheme is properly configured and that the backups are reliable.

upvoted 1 times

🗨️ 👤 **kloug** 2 years, 3 months ago

cccccccccccccc

upvoted 1 times

🗨️ 👤 **Obi_Wan_Jacoby** 2 years, 5 months ago

Selected Answer: D

I am going with Answer D. Sure, a test (C) is also needed, and may very well work, but you need to verify media integrity and confirm it is not in need of replacement. The very first thing in the All-In-One Server+ book "after" mentioning picking out a scheme (backup solution), is to verify media integrity.

upvoted 2 times

🗨️ 👤 **szl0144** 2 years, 9 months ago

D seems correct.

upvoted 3 times

A systems administrator is performing maintenance on 12 Windows servers that are in different racks at a large datacenter. Which of the following would allow the administrator to perform maintenance on all 12 servers without having to physically be at each server? (Choose two.)

- A. Remote desktop
- B. IP KVM
- C. A console connection
- D. A virtual administration console
- E. Remote drive access
- F. A crash cart

Suggested Answer: AB

Reference:

<https://www.blackbox.be/en-be/page/27559/Resources/Technical-Resources/Black-Box-Explains/kvm/Benefits-of-using-KVM-over-IP>

1. Remotely Access Computers and Software

Take control of all remote computers, servers and virtual machines on your LAN or WAN by connecting your KVM switches over your existing IP network. KVM over IP combines the advantages of remote access software with the benefits of KVM switching technology. Like most KVM switches, KVMoIP products don't require any software to be loaded on the host computers. They interface directly with the keyboard, monitor, and mouse connectors of the host computer or KVM switch. The KVM over IP switch digitises the incoming video signal and processes it into digital data that is communicated to a remote client computer over a LAN/WAN or the public Internet.

  **PaytoPlay** 10 months, 1 week ago

I would go B & D here. D, because I use this everyday at work managing 20 servers with vSphere. B, because we use this as well on some of our older equipment. I also use A as well but less often.

upvoted 3 times

  **Julienzen** 2 years, 1 month ago

A,B is correct answers.

upvoted 1 times

A server administrator is experiencing difficulty configuring MySQL on a Linux server. The administrator issues the `getenforce` command and receives the following output:

```
># Enforcing
```

Which of the following commands should the administrator issue to configure MySQL successfully?

- A. `setenforce 0`
- B. `setenforce permissive`
- C. `setenforce 1`
- D. `setenforce disabled`

Suggested Answer: A

Reference:

<https://blogs.oracle.com/mysql/selinux-and-mysql-v2>

🗨️ 👤 **RBL23168** 7 months, 1 week ago

The `getenforce` command is used to check the current status of SELinux (Security-Enhanced Linux) on a system. The output "Enforcing" indicates that SELinux is currently in enforcing mode, which may restrict certain operations.

To configure MySQL successfully, the administrator can temporarily set SELinux to permissive mode, which allows policy violations to be logged but not enforced. The correct command for this is:

B. `setenforce permissive`

This command will set SELinux to permissive mode, allowing the administrator to identify and address any SELinux policy violations related to MySQL without enforcing strict policies. After configuring MySQL, it's generally a good practice to set SELinux back to enforcing mode for enhanced security, using the command:

C. `setenforce 1`

This will set SELinux back to enforcing mode

upvoted 1 times

🗨️ 👤 **RBL23168** 7 months, 1 week ago

A. `setenforce 0`. This command sets SELinux to "Permissive" mode. In Permissive mode, SELinux does not deny access but logs AVC (Access Vector Cache) messages. It's useful for troubleshooting and understanding how SELinux would behave without actually enforcing any denials.

upvoted 2 times

🗨️ 👤 **Nasiim** 11 months, 1 week ago

The answer is A

Setenforce 0: to change to permissive mode

Setenforce 1: to change to enforcing mode

upvoted 1 times

🗨️ 👤 **jjwelch00** 1 year, 1 month ago

"setenforce 0" sets SELinux to "disabled" mode, which completely turns off SELinux enforcement.

"setenforce permissive" sets SELinux to "permissive" mode, which allows SELinux to log policy violations but does not block any actions.

"setenforce 1" sets the SELinux enforcement to "enforcing" mode, which means SELinux policy rules are actively enforced.

"setenforce disabled" sets the SELinux enforcement to "disabled" mode, which means SELinux policy rules are not enforced at all, and the system falls back to standard Unix permissions.

upvoted 1 times

Which of the following backup types only records changes to the data blocks on a virtual machine?

- A. Differential
- B. Snapshot
- C. Incremental
- D. Synthetic full

Suggested Answer: C

Reference:

<https://searchdatabackup.techtarget.com/definition/incremental-backup>

🗨️ **ompk** 9 months ago

Chatgpt and other websites are choosing C: Snapshots
upvoted 2 times

🗨️ **ompk** 8 months, 1 week ago

My bad guys, its Incremental
upvoted 1 times

🗨️ **eino** 2 years, 6 months ago

I don't like how this is written, because it could be incremental or differential, in my thinking. Yes, incremental will only copy the changed data blocks, but then it ALSO resets the archive bit.... Where differential will only backup what has changed since the last time the bit was reset.
upvoted 1 times

🗨️ **Pongsathorn** 2 years, 7 months ago

I think a keyword is "only records changes to the data blocks". The provided answer could be correct.
upvoted 1 times

🗨️ **Dion79** 3 years, 2 months ago

https://helpcenter.veeam.com/docs/backup/hyperv/changed_block_tracking.html?ver=110
upvoted 1 times

🗨️ **Dion79** 3 years ago

<https://www.ibm.com/docs/en/spfve/8.1.2?topic=machines-vm-backups-resilient-change-tracking-rct>

Tricky question.. Provided answer maybe correct I was thinking snapshot but question seem smore geeared to incremental.
upvoted 2 times

🗨️ **[Removed]** 10 months, 2 weeks ago


Snapshots aren't backups, so the answer is definitely Incremental.
upvoted 1 times

Which of the following server types would benefit MOST from the use of a load balancer?

- A. DNS server
- B. File server
- C. DHCP server
- D. Web server

Suggested Answer: *D*

Community vote distribution



Biancoega10 10 months, 1 week ago

Selected Answer: D

I think it's D.

upvoted 2 times

RBL23168 1 year, 4 months ago

Selected Answer: D

Clearly D. LBs distribute network traffic for performance.

upvoted 3 times

AzadOB 1 year, 5 months ago

Selected Answer: A

file server

upvoted 1 times

A company uses a hot-site, disaster-recovery model. Which of the following types of data replication is required?

- A. Asynchronous
- B. Incremental
- C. Application consistent
- D. Constant

Suggested Answer: D

Community vote distribution

D (71%)

A (29%)

🗳️ 👤 **RBL23168** Highly Voted 1 year, 7 months ago

There is no such thing as Constant Replication, it's proper name is Synchronous replication.

Synchronous replication ensures that data is written to primary and alternate locations without delay; this results in an up-to-date mirror copy of data between a primary and a hot site and is often done in the background automatically. Asynchronous replication includes a slight delay before data is written to alternate sites; as a result, this is less expensive than synchronous solutions, but it can cause problems with applications depending upon database consistency. A B C are all not valid options for me so i'd go with D. Typically vague CompTIA question.

upvoted 5 times

🗳️ 👤 **FreePrivacy** Most Recent 11 months, 1 week ago

Explanation:

Constant Replication

The primary system replicates data changed data blocks continually. Constant replication is also referred to as "continuous replication." The replication process occurs in the background, permitting users to access the data without interruption.

Constant replication is different than regular backups, where files must be closed to be duplicated.

upvoted 1 times

🗳️ 👤 **Sweety_Certified7** 1 year, 5 months ago

Selected Answer: D

First i thought A should be the answer but: Asynchronous data replication is typically associated with a warm-site disaster recovery model, not a hot-site model. In a hot-site disaster recovery setup, synchronous (or, in this case; constant) replication is more commonly used, where the secondary site is fully operational and mirrors data in real-time from the primary site. so D is correct.

upvoted 2 times

🗳️ 👤 **jjwelch00** 2 years, 1 month ago

Answer is D.

upvoted 2 times

🗳️ 👤 **Katlegobogosi** 2 years, 4 months ago

Selected Answer: A

A. Asynchronous.

Asynchronous data replication is typically used in a hot-site disaster recovery model. This is because asynchronous replication does not require an immediate update to the remote site, allowing for a slight delay in replication. This delay is usually acceptable in a hot-site model, where the remote site is already configured and ready to take over in case of a disaster. Option A is the correct answer.

upvoted 3 times

🗳️ 👤 **Obi_Wan_Jacoby** 2 years, 5 months ago

Selected Answer: D

Answer is D. Answer A technically runs the risk of data loss.

upvoted 1 times

🗄️ 👤 **eino** 2 years, 6 months ago

Selected Answer: D

Constant.

A hot site requires the same data that's at your primary site, requiring constant replication.

upvoted 1 times

🗄️ 👤 **Pongsathorn** 2 years, 8 months ago

Selected Answer: D

Constant

When constant replication is in use, each time a change is made (data addition, deletion, change, etc.) at the primary site, the same change is written to the secondary site. This results in the secondary site being constantly up-to-date and is generally an expensive option.

CompTIA Server+ Study Guide Exam SK0-005 Chapter 9 Disaster Recovery

upvoted 4 times

🗄️ 👤 **nixonbii** 2 years, 8 months ago

Does CompTia use its own exam guides to write the tests? According to their book, chapter 8 page 329, the two types of replication used for a hot site are synchronous and asynchronous. There is no mention anywhere in the text about "constant" replication. I hope they are not sourcing test material from outside of the book that they produced for the purpose of learning these topics.

upvoted 2 times

🗄️ 👤 **Obi_Wan_Jacoby** 2 years, 5 months ago

Read page 328 at the very top. It helped me

upvoted 1 times

🗄️ 👤 **jagoichi** 2 years, 9 months ago

Constant Replication

The primary system replicates data changed data blocks continually. Constant replication is also referred to as "continuous replication." The replication process occurs in the background, permitting users to access the data without interruption. Constant replication is different than regular backups, where files must be closed to be duplicated.

upvoted 1 times

🗄️ 👤 **Timock** 2 years, 9 months ago

Hot sites are actively updated asynchronously. This occurs in real time, and provides a near-mirror image of your production site on your target systems. Standby latencies for hot sites are typically only milliseconds in length, resulting in little to no downtime during failover.

Asynchronous replication is also commonly used in virtual machines (VMs): Some hypervisors include asynchronous replication to allow entire VMs to be replicated to a remote location so the VM can fail over to that location in the event of a disaster. Another common use case for asynchronous replication is storage snapshots for continuous data protection.

<https://prototechnologies.com/blog/disaster-recovery/recovery-site-hot-warm-or-cold/>

<https://blog.purestorage.com/purely-informational/synchronous-replication-vs-asynchronous-replication/#>

upvoted 3 times

🗄️ 👤 **szl0144** 2 years, 9 months ago

Selected Answer: D

D must be the answer

upvoted 1 times

🗄️ 👤 **PEsty93** 3 years, 2 months ago

Selected Answer: A

"Constant" isn't a replication method. The methods are synchronous or asynchronous.

Asynchronous replicates the data at a scheduled time and reduces bandwidth, compared to synchronous which replicates data instantly.

Synchronous would be best if you have two live sites. Asynchronous would be used for a hot-site which would only be used for a failover.



upvoted 1 times

🗄️ 👤 **Dion79** 3 years, 1 month ago

It is a replication method...., according to CompTIA. It's a trick questions and I'd probably still pick Constant.. Page 249 Topic 11c (Managing Service and Data Availability - The Office CompTIA Server+ Study Guide Exam Sk0-005)

Constant, Periodic, Asynchronous, Synchronous, Application Consistent, File Locking, Mirroring, BiDi...



upvoted 2 times

  **Dion79** 3 years, 2 months ago

Selected Answer: D

Constant replication over a high-speed network link can ensure fast data copies and data consistency across locations.

upvoted 1 times

  **Dion79** 3 years, 1 month ago

The primary system replicates data changed data blocks continually. Constant replication is also referred to as "continuous replication." The replication process occurs in the background, permitting users to access the data without interruption. Constant replication is different than regular backups, where files must be closed to be duplicated.

asynchronous replication, data is sent to the primary storage (a database, for example). Confirmation that the data was written is sent to the client machine. The data is then replicated to the second database server, and it responds to the first database server with a confirmation. The problem arises if the primary server fails. The secondary may not actually have the replicated data, but the client computer believes that it does. The data is then lost.

Reference: The Official CompTIA Server+ Study Guide


upvoted 5 times

A technician is unable to access a server's package repository internally or externally. Which of the following are the MOST likely reasons? (Choose two.)

- A. The server has an architecture mismatch
- B. The system time is not synchronized
- C. The technician does not have sufficient privileges
- D. The external firewall is blocking access
- E. The default gateway is incorrect
- F. The local system log file is full

Suggested Answer: CD

Community vote distribution



Katlegobogosi Highly Voted 2 years, 4 months ago

Selected Answer: CD

- C. The technician does not have sufficient privileges to access the package repository.
- D. The external firewall is blocking access to the package repository.

Therefore, options C and D are the correct answers.

Option A (architecture mismatch) may cause issues with software installation or updates, but it is not related to the server's package repository access.

Option B (system time not synchronized) may cause authentication issues but is unlikely to affect package repository access.

Option E (incorrect default gateway) may affect network connectivity, but it is not directly related to package repository access.

Option F (local system log file full) may affect logging and monitoring but is unlikely to affect package repository access.

upvoted 5 times

Replicant 2 years, 3 months ago

The question said INTERNALLY and externally so D is incorrect about the external firewall being issue.

upvoted 1 times

ompk 9 months ago

Hello Dear Replicant, read the question again as it says INTERNALLY or EXTERNALLY. so it means both

upvoted 1 times

Kabbath1986 Most Recent 2 months, 1 week ago

Selected Answer: DE

- D. The external firewall is blocking access
- E. The default gateway is incorrect

Here's why:

The technician is unable to access the package repository, internally or externally – meaning the issue could be related to network access or routing problems, not necessarily something like permissions or logs.

upvoted 1 times

ompk 9 months ago

Hello

For people choosing D & E, both answers are correct for external access issues but what about the internal access issue? out of A,B,C and F, C is the most appropriate answer for not having access INTERNALLY to a repository for a technician. In this context I think C and D are the right answers where D emphasizes more on the network side rather accessing files in the server.

upvoted 1 times

kx7tg4xu 10 months, 1 week ago

Selected Answer: CD

- C. The technician does not have sufficient privileges

If the technician does not have sufficient privileges to access the repository, the internal repository cannot be accessed. For internal access, this is

most likely due to the technician's lack of authority.

D. The external firewall is blocking access

If access to the external repository is not possible, it is most likely that a firewall is blocking traffic.

upvoted 1 times

  **Kraken84** 1 year, 8 months ago

The most likely reasons a technician is unable to access a server's package repository both internally and externally are:

D. The external firewall is blocking access - Firewalls can block incoming and outgoing traffic. If the firewall is configured to block access to the package repository, the technician won't be able to access it.

E. The default gateway is incorrect - The default gateway is responsible for routing traffic outside the local network. If it's incorrectly configured, the server might not be able to reach external resources, including external package repositories.

While some of the other options might cause issues with specific operations on a server, they are less likely to be the direct cause of inability to access a package repository both internally and externally.


upvoted 2 times

  **Dingos** 2 years, 1 month ago

Selected Answer: CE

from my experience C it is sure right from other gateway misconfiguration can happen if there is some VLAN change.

upvoted 1 times

  **kloug** 2 years, 3 months ago

dddddddddddddddddd,eeeeeeeeeeeeeeeeee



upvoted 1 times

  **Pongsathorn** 2 years, 8 months ago

Selected Answer: AC

A, C should be the answer.

upvoted 3 times

  **Dion79** 3 years, 4 months ago

I also think C could be a valid answer as well... if not having the correct permissions can also cause this.

upvoted 2 times

A server administrator was asked to build a storage array with the highest possible capacity. Which of the following RAID levels should the administrator choose?

- A. RAID 0
- B. RAID 1
- C. RAID 5
- D. RAID 6



Suggested Answer: A

Reference:

<https://www.steadfast.net/blog/almost-everything-you-need-know-about-raid>

  **Greedy1985** 7 months, 3 weeks ago



I think that would be RAID 5. Source is here <https://www.trentonsystems.com/blog/raid-levels-0-1-5-6-10-raid-types>
upvoted 1 times

  **Pongsathorn** 7 months, 3 weeks ago

"highest possible capacity"

RAID 0 offers highest capacity compared to the other RAID's.

upvoted 3 times

  **Dion79** 1 year, 4 months ago

Answer looks correct...

<https://www.computerweekly.com/answer/Choosing-a-RAID-level-depends-on-capacity-needs-and-data-criticality>

upvoted 1 times

A technician needs to deploy an operating system that would optimize server resources. Which of the following server installation methods would BEST meet this requirement?

- A. Full
- B. Bare metal
- C. Core
- D. GUI

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **RBL23168** 7 months, 1 week ago

C.. easiest question on the exam. If you don't know this then idk... . The Core installation is a minimalistic installation option that includes only the essential components needed to run the operating system

upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 1 year, 5 months ago

Selected Answer: C

I believe C is correct. I researched this quite a bit. Then in the book "All-In-One CompTia Server+ Certification" Exam Guide, on page 94 I found what they (CompTia) are looking for. They referred to Server Installation "Methods" and listed "Core" (as well as GUI). Obviously, Core will take less resources. "Bare Metal" is used for better performance and having direct access to the hardware (as you do not share it with other clients). Also (fyi) a Bare Metal Hypervisor is a Bare Metal server but with a Hypervisor installed which also restricts some access to direct hardware (type1 vs type2). Info on that can be found here "<https://phoenixnap.com/blog/what-is-bare-metal-server>" I do not believe "Bare Metal" is considered a "method" for installing an OS.

upvoted 2 times

🗳️ 👤 **Pongsathorn** 1 year, 7 months ago

Selected Answer: C

Windows Servers

Windows servers have traditionally been managed via a GUI. Beginning with Windows Server 2008, Microsoft has offered the Server Core installation option. Server Core is a command line-only OS. It has a much smaller hardware footprint on the server and can provide significant performance benefits.

Ref. The Official CompTIA Server+ Study Guide (Exam SK0-005) page 129.

upvoted 2 times

🗳️ 👤 **Fineb** 1 year, 9 months ago

Core is the correct answer, if the question was asked about hypervisor type then we can say it's bare metal. Here, the question was about the OS itself and Core is more rapid and less resources consumption because no GUI, just like a CLI and it's also secure that normal GUI OS

upvoted 1 times

🗳️ 👤 **Fineb** 1 year, 9 months ago

** secure than GUI OS

upvoted 1 times

🗳️ 👤 **Landoski** 1 year, 9 months ago

B. Bare metal has access to all the resources

upvoted 1 times

🗳️ 👤 **jagoichi** 1 year, 9 months ago

Answer is B bare mental

Bare Metal OS installed directly on the server hardware

Traditional installation option

Resources not shared among multiple installations

Provides the security, reliability, and performance required by

some industries

May be more costly than virtualization

upvoted 2 times

A company's IDS has identified outbound traffic from one of the web servers coming over port 389 to an outside address. This server only hosts websites. The company's SOC administrator has asked a technician to harden this server. Which of the following would be the BEST way to complete this request?

- A. Disable port 389 on the server
- B. Move traffic from port 389 to port 443
- C. Move traffic from port 389 to port 637
- D. Enable port 389 for web traffic

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ **broman** 9 months, 3 weeks ago

I think the answer C was a typo 636 is the secure port for LDAP and would be correct, however given the options, A is correct
upvoted 1 times

🗳️ **AzadOB** 9 months, 3 weeks ago

Selected Answer: A

Port 389 is commonly associated with LDAP (Lightweight Directory Access Protocol), which is used for directory services. Since the web server should not be conducting LDAP-related activities, it's concerning that outbound traffic is being observed on this port. To harden the server and ensure that it's not misused for unintended purposes, the best approach would be to:

A. Disable port 389 on the server.
upvoted 2 times

🗳️ **kloug** 1 year, 9 months ago

aaaaaaaaaaaa

upvoted 1 times

🗳️ **Pongsathorn** 2 years, 1 month ago

Selected Answer: A

OS Hardening

Hardening the server should start with hardening the operating system. This involves a series of steps that should result in a server that offers a minimum of attack points to a hacker. Let's look at six steps that can lead to this result.

Disable Unused Services/Close Unneeded Ports

Any services that are not required on the server should be disabled. Only those required for the server to perform its role in the network should be left on. The easiest way to do this is to install a host firewall on the system and adopt a "disable by default" policy with respect to services by closing the port used for the service. Then manually enable any you need.

Ref. CompTIA Server+ Study Guide: Exam SK0-005 Troy McMillan

upvoted 1 times

🗳️ **nixonbii** 2 years, 2 months ago

Selected Answer: A

We all want to give the test the answer it wants but if you discover a server handling traffic that is outside of its scope of operation, you need to shut down the offending port ASAP. Complete security scans and try to find out what the nature of that traffic was.

upvoted 1 times

🗳️ **Timock** 2 years, 3 months ago

Selected Answer: A

This is solely a web server so should be ports for HTTPS. There should not be any traffic leaving the server on 389 externally and its LDAP traffic so definitely a concern. Port 389 can be disabled. Not to mention LDAPS would be 636 not 637.

http://ports.my-addr.com/tcp_port-udp_port-application-and-description.php?port=637



upvoted 1 times

  **TheITStudent** 2 years, 4 months ago

Selected Answer: A

For this, sense this only needs web traffic hosted, only port 80 & 443 need to be open, port 389 can be closed... I would just disable it. The better answer would be a more specific firewall rule set, but given the options, I would choose A.

upvoted 4 times

  **Dion79** 2 years, 7 months ago

A or C??

upvoted 1 times

Which of the following would be BEST to help protect an organization against social engineering?

- A. More complex passwords
- B. Recurring training and support
- C. Single sign-on
- D. An updated code of conduct to enforce social media

Suggested Answer: B

Community vote distribution

B (100%)

  **PEsty93**  1 year, 8 months ago

Selected Answer: B

Training. How would enforcing social media help anything?
upvoted 9 times

  **Dion79**  1 year, 9 months ago

User awareness and training is the number one defense against most security threats. Chapter 6 discussed social engineering and user awareness regarding network malware attacks.
Reference: All in-one CompTIA Server+ Certification Exam Guide SK0-005
upvoted 7 times

  **Obi_Wan_Jacoby**  11 months, 2 weeks ago

Selected Answer: B

B is correct answer
upvoted 1 times

A technician is connecting a server's secondary NIC to a separate network. The technician connects the cable to the switch but then does not see any link lights on the NIC. The technician confirms there is nothing wrong on the network or with the physical connection. Which of the following should the technician perform NEXT?

- A. Restart the server
- B. Configure the network on the server
- C. Enable the port on the server
- D. Check the DHCP configuration

Suggested Answer: C

Community vote distribution

C (67%)

B (33%)

🗳️ 👤 **Ronn_Burgandy** 9 months, 1 week ago

Selected Answer: C

No link lights and the network configuration has been verified would suggest something is disabled with the port. If it was only a configuration issue you should still see some link light activity.

upvoted 1 times

🗳️ 👤 **PaytoPlay** 10 months, 1 week ago

B, if you're going through the trouble of connecting the second NIC you should probably configure it first before connecting it to the network.

upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 2 years, 5 months ago

Selected Answer: C

C looks correct. If you disable your NIC or uninstall it (or if it is corrupted), I believe the light will go out (if you had it on prior). Restarting (A) is part of troubleshooting, B&C would not cause the light to not come on.

upvoted 3 times

🗳️ 👤 **eino** 2 years, 6 months ago

Selected Answer: B

Nothing in the question states that the technician has configured the NIC for the "new" network it is getting setup for.

upvoted 2 times

🗳️ 👤 **TheITStudent** 2 years, 10 months ago

Vague question... as always! My first thoughts are that this would need to be configured... but since "everything on the network is fine"... whatever that means... I would guess that maybe there is port security and to turn that on... also an inference... this question stinks.

upvoted 2 times

Which of the following would MOST likely be part of the user authentication process when implementing SAML across multiple applications?

- A. SSO
- B. LDAP
- C. TACACS
- D. MFA

Suggested Answer: A

Reference:

<https://www.onelogin.com/learn/how-single-sign-on-works>

Community vote distribution

A (100%)

🗨️ **weat** 8 months, 1 week ago

I think D - the user has to authenticate, usually with 2 MFA, to the authentication server via the SSO platform, which then informs the particular service requested, that the user is authenticated, and the service gives the user access. "Security Assertion Markup Language (SAML) is a protocol that enables SSO " - it's not what the user does. <https://cloudinfrastructureservices.co.uk/saml-vs-sso-whats-the-difference/>

upvoted 1 times

🗨️ **Obi_Wan_Jacoby** 11 months, 2 weeks ago

Selected Answer: A

A is correct

upvoted 2 times

🗨️ **magi420** 4 months, 2 weeks ago

thank you obi_wan_jacoby your our only hope. lol love the username.

upvoted 1 times

A server administrator needs to check remotely for unnecessary running services across 12 servers. Which of the following tools should the administrator use?

- A. DLP
- B. A port scanner
- C. Anti-malware
- D. A sniffer


Suggested Answer: B

Reference:

<https://www.getsafeonline.org/business/articles/unnecessary-services/>

Community vote distribution

B (100%)

  **Timock** 9 months, 2 weeks ago

Selected Answer: B

Sniffing is the term generally used for traffic monitoring within a network, while port scanning is used to find out information about a remote network. Both sniffing and port scanning have the same objective—to find system vulnerabilities—but they take different approaches.

<https://www.informit.com/articles/article.aspx?p=31964#>

upvoted 2 times

A company is building a new datacenter next to a busy parking lot. Which of the following is the BEST strategy to ensure wayward vehicle traffic does not interfere with datacenter operations?

- A. Install security cameras
- B. Utilize security guards
- C. Install bollards
- D. Install a mantrap

Correct Answer: C


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

A technician has been asked to check on a SAN. Upon arrival, the technician notices the red LED indicator shows a disk has failed. Which of the following should the technician do NEXT, given the disk is hot swappable?

- A. Stop sharing the volume
- B. Replace the disk
- C. Shut down the SAN
- D. Stop all connections to the volume

Suggested Answer: B

Community vote distribution



🗳️ 👤 **Hitemup** 7 months, 3 weeks ago

Just replace the disk since it's hot swappable
upvoted 2 times

🗳️ 👤 **Riseofashes** 8 months, 2 weeks ago

Selected Answer: B

Just to balance Obi_wans realization.
It should be B, the disk is hot-swappable, there's no reason to disrupt any operations.
upvoted 2 times

🗳️ 👤 **azre_certified1111** 9 months, 3 weeks ago

Selected Answer: B

Easily replace the disk because its hot swappable.
upvoted 2 times

🗳️ 👤 **Obi_Wan_Jacoby** 11 months, 2 weeks ago

Selected Answer: D

I believe D is correct. You would not want to stop the share (how many shares/printers have been broken this way?). I would stop the connections first, then shut it down (as it cannot be hot swapped in this case) then replace the disk.
upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 11 months, 2 weeks ago


Well, I should re-read prior to submitting, LOL.. I imagined reading it is NOT hot swappable! When in-fact it says it IS.. My bad, replace disk is correct
upvoted 4 times

Network connectivity to a server was lost when it was pulled from the rack during maintenance. Which of the following should the server administrator use to prevent this situation in the future?

- A. Cable management
- B. Rail kits
- C. A wireless connection
- D. A power distribution unit

Suggested Answer: A

Community vote distribution



🗳️ 👤 **Riseofashes** Highly Voted 1 year, 8 months ago

Selected Answer: A

From my experience, rail-kits refers to the two steel bars that you use to install then move the server in and out, but it doesn't have anything to do with the cables.

A cable management arm is specifically what you would use if you want to extend the server out of the rack without removing cables. So I choose "Cable management".

upvoted 6 times

🗳️ 👤 **Kraken84** 1 year, 2 months ago

If only that word ARM was in the answer, would have been a dead giveaway

upvoted 1 times

🗳️ 👤 **cj207800** Most Recent 5 months ago

Selected Answer: B

Rail kits are designed to allow servers to be easily slid in and out of racks while maintaining all necessary connections, including network cables. They provide a stable mounting solution that keeps the server securely attached to the rack even when extended for maintenance.

upvoted 1 times

🗳️ 👤 **Chiaretta** 11 months ago

Selected Answer: B

rails kit is the answer

upvoted 1 times

🗳️ 👤 **kloug** 1 year, 9 months ago

aaaaaaaaaaaaaa

upvoted 1 times

🗳️ 👤 **azre_certified1111** 1 year, 9 months ago

Selected Answer: B

I believe rail kits is correct. The whole point of rail kits its to ensure that server equipment are properly seated and won't be accidentally pulled out during maintenance or other activities

upvoted 1 times

🗳️ 👤 **Jfrican** 1 year, 9 months ago

A rail kit would include cable management and would prevent from being pulled out.

I guess cable management would be the cheapest though.

It is a vague question.

upvoted 1 times

Which of the following access control methodologies can be described BEST as allowing a user the least access based on the jobs the user needs to perform?

- A. Scope-based
- B. Role-based
- C. Location-based
- D. Rule-based

Suggested Answer: B

Community vote distribution

B (100%)

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

Selected Answer: B

"based on the jobs the user needs to perform" I.e. based on their role in the company.
upvoted 2 times


A datacenter technician is attempting to troubleshoot a server that keeps crashing. The server runs normally for approximately five minutes, but then it crashes.

After restoring the server to operation, the same cycle repeats. The technician confirms none of the configurations have changed, and the load on the server is steady from power-on until the crash. Which of the following will MOST likely resolve the issue?

- A. Reseating any expansion cards in the server
- B. Replacing the failing hard drive
- C. Reinstalling the heat sink with new thermal paste
- D. Restoring the server from the latest full backup

Suggested Answer: C

Community vote distribution



🗨️ 👤 **Ronn_Burgandy** 9 months, 1 week ago

Selected Answer: C

C would be the only answer to explain why there is an issue after about 5 minutes. Just long enough for the CPU to overheat. There are over reasons (ie rouge processes) why an issue could happen over a period of time but C is only one out of this group of answers. A. doesn't make sense how would expansion card become loose after running the server for 5 minutes? Reseating the expansion card would make more sense if the server couldn't see the card at all.

upvoted 1 times

🗨️ 👤 **kx7tg4xu** 10 months, 1 week ago

Selected Answer: C

C is the correct answer.

Problems that occur over time often indicate thermal problems.

upvoted 1 times

🗨️ 👤 **Rookert** 10 months, 1 week ago

Selected Answer: A

Wouldn't a server upon reaching Tj (maximum temperature) just simply circuit itself to protect it - that means the power CUTS OFF.

A system crash is not equal to a sudden shutdown. I would go for the reseal, A..

upvoted 1 times

🗨️ 👤 **Rookert** 10 months, 1 week ago

Wouldn't a server upon reaching Tj (maximum temperature) just simply circuit itself to protect it - that means the power CUTS OFF.

A system crash is not equal to a sudden shutdown. I would go for the reseal, A..

upvoted 1 times

🗨️ 👤 **Riseofashes** 2 years, 2 months ago

Selected Answer: C

A server suddenly turning off without any power issues typically indicates overheating.

CPUs overheating can very often cause servers to just die after being on for 2-10 minutes.

upvoted 1 times

A server administrator is exporting Windows system files before patching and saving them to the following location:

\\server1\ITDept\

Which of the following is a storage protocol that the administrator is MOST likely using to save this data?

- A. eSATA
- B. FCoE
- C. CIFS
- D. SAS

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Kraken84** 8 months, 2 weeks ago

Selected Answer: C

C. CIFS

The provided location (\\server1\ITDept\) is a UNC (Universal Naming Convention) path, which is commonly used with the CIFS (Common Internet File System) protocol. CIFS is a network file-sharing protocol and is the standard way that Windows computers share files with other devices on a network.

upvoted 3 times

🗳️ 👤 **Obi_Wan_Jacoby** 1 year, 5 months ago

C looks correct to me (guess that would be SMB2 and 3 now?)

upvoted 1 times

🗳️ 👤 **[Removed]** 1 year, 9 months ago

How can the answer be CIFS (SMB)? Look at the back slashes in the question's path for the network drive in the question and they do NOT indicate that it is a Linux type of network path share, but a Windows type of network path share, i.e., Windows uses backslashes whereas Linux path to network drives uses forward slashes, e.g., see: <https://www.minitool.com/lib/unc-path.html> Not sure how they can cite the answer as CIFS given the information provided in the question.

upvoted 1 times

🗳️ 👤 **Dion79** 2 years, 4 months ago

Common Internet File System (CIFS) or Server Message Blocks (SMB)

Network access to shared folders

TCP 137, 139, UDP 137, 139

The Common Internet File System (CIFS) is a dialect or an implementation of the SMBs protocol. Most references will be to SMB, and that is the proper term by today's standards. One of the most common servers on the network are file servers. These devices are designed to store large numbers of files. Workstations access to these files from across the network. Application layer protocols manage file access. The two most common protocols are the Network File System (NFS) and Server Message Blocks (SMB).

upvoted 2 times

🗳️ 👤 **i_bird** 1 year, 10 months ago

TCP: 139 and 445, 445 being CIFS over IPv6

UDP: 137 and 138

upvoted 1 times

A server technician has received reports of database update errors. The technician checks the server logs and determines the database is experiencing synchronization errors. To attempt to correct the errors, the technician should FIRST ensure:

- A. the correct firewall zone is active
- B. the latest firmware was applied
- C. NTP is running on the database system
- D. the correct dependencies are installed

Suggested Answer: C

Community vote distribution

C (100%)

🗉 👤 **Ronn_Burgandy** 9 months, 1 week ago

Selected Answer: C

Keyword is Synchronization. Network Time Protocol is crucial for servers and data consistency. All other answers may affect connectivity or functionality but not with syncing.

upvoted 1 times

A technician is connecting a Linux server to a share on a NAS. Which of the following is the MOST appropriate native protocol to use for this task?

- A. CIFS
- B. FTP
- C. SFTP
- D. NFS

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **shanebrown** 9 months, 2 weeks ago

Selected Answer: D

CIFS is optimized for windows but can be unused on linux, while NFS is optimized for linux but can be used on windows.
upvoted 1 times

🗳️ 👤 **Evanj51** 2 years, 7 months ago

Selected Answer: D

Linux uses NFS
upvoted 2 times

🗳️ 👤 **momoci** 2 years, 7 months ago

Selected Answer: D

Please mr admin. change answer to D.
NFS (Network File System)
CIFS (Common Internet File System)
upvoted 1 times

🗳️ 👤 **Pongsathorn** 2 years, 8 months ago

Selected Answer: D

The answer is D.

Network File System (NFS)

NFS is one of the protocols that can be used to connect to and share data on a NAS device. Network File System (NFS) is a client/server file-sharing protocol used in Unix/Linux. Version 4 is the most current version of NFS. It operates over TCP port 2049. Secure NFS (SNFS) offers confidentiality using Digital Encryption Standard (DES).

Common Internet File System (CIFS)

Server Message Block (SMB) is an application layer protocol used to provide shared access to resources. The Common Internet File System (CIFS) protocol is a dialect of SMB. It is primarily used in Windows systems. The latest version is 3.1.1, which was released to support Windows 10 and Windows Server 2016. It operates as a client-server application. It uses port 445.

CompTIA Server+ Study Guide: Exam SK0-005 Chapter 4 Storage Technologies and Asset Management

upvoted 3 times

🗳️ 👤 **Timock** 2 years, 9 months ago

Selected Answer: D

Network File System (NFS): NFS is a network that was introduced by Sun Microsystems and is used by Unix or Linux-based operating systems and stands for Network File System. This is a network that is used for giving remote access capabilities to the applications. Remote access enables the user to edit or even take a closer look at his computer by using another computer. Old files can be repaired even when the user is at a distance from his computer. This protocol gives devices the functionality to modify the data over a network.

<https://www.geeksforgeeks.org/difference-between-nfs-and-cifs/#>

upvoted 1 times

🗳️ 👤 **Timock** 2 years, 9 months ago

Even the explanation states NFA for Linux and CIFS for Windows... but for some odd reason CIFS has been chosen as the answer here.
upvoted 1 times

A server in a remote datacenter is no longer responsive. Which of the following is the BEST solution to investigate this failure?

- A. Remote desktop
- B. Access via a crash cart
- C. Out-of-band management
- D. A Secure Shell connection

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Dion79** Highly Voted 2 years, 8 months ago

Selected Answer: C

Out-of-band administration provides hardware-level remote access to a host without relying on the OS software running. Common solutions include Dell's iDRAC and HP's iLO; the server hardware must support this type of remote administration, and it must be configured with IP settings.

upvoted 12 times

🗳️ 👤 **ahmsha** Most Recent 9 months, 2 weeks ago

the correct answer is selected D Secure Shell (SSH) is a network protocol that provides a secure way for users, particularly system administrators, to access a computer over an unsecured network

upvoted 1 times

🗳️ 👤 **Pongsathorn** 2 years, 2 months ago

Selected Answer: C

Out-of-Band Management

Out-of-band (OOB) management refers to any method of managing the server that does not use the network. This provides some advantages, among them:

It offers a solution when the network is down or the device is inaccessible.

It manages devices with no power and remotely reboots devices that have been crashed, turned off, hibernating, or in sleep mode.

There are various ways to connect to a server without using the network, and in this section you'll learn about some of these methods.

upvoted 1 times

🗳️ 👤 **Timock** 2 years, 3 months ago

Selected Answer: C

SSH cannot be the answer as it would require a functioning O/S on the other side. Out-of-band solutions is the correct answer here.

upvoted 1 times

🗳️ 👤 **haazybanj** 2 years, 9 months ago

C seems more appropriate

upvoted 4 times

🗳️ 👤 **Pieman125** 3 years, 2 months ago

Answer should be C.

upvoted 4 times

A server is reporting a hard drive S.M.A.R.T. error. When a technician checks on the drive, however, it appears that all drives in the server are functioning normally. Which of the following is the reason for this issue?

- A. A S.M.A.R.T. error is a predictive failure notice. The drive will fail in the near future and should be replaced at the next earliest time possible
- B. A S.M.A.R.T. error is a write operation error. It has detected that the write sent to the drive was incorrectly formatted and has requested a retransmission of the write from the controller
- C. A S.M.A.R.T. error is simply a bad sector. The drive has marked the sector as bad and will continue to function properly
- D. A S.M.A.R.T. error is an ECC error. Due to error checking and correcting, the drive has corrected the missing bit and completed the write operation correctly.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **Dion79** Highly Voted 3 years, 2 months ago

Selected Answer: A

correct answer is A

upvoted 6 times

🗳️ 👤 **haazybanj** Highly Voted 3 years, 3 months ago

A is correct

upvoted 5 times

🗳️ 👤 **shanebrown** Most Recent 9 months, 2 weeks ago

Selected Answer: A

SMART indicates that the drive has a potential problem

upvoted 1 times

🗳️ 👤 **Pongsathorn** 2 years, 8 months ago

Selected Answer: A

Correct answer is A"

Predictive Failures

Some issues will issue a warning that a failure is in the future. An example of such a device is a hard drive that has support for Self-Monitoring, Analysis and Reporting Technology (SMART). This technology is built into most drives to determine if the device is still physically healthy or failing due to hardware issues.

upvoted 2 times

🗳️ 👤 **szl0144** 2 years, 9 months ago

Selected Answer: A

A is correct

upvoted 1 times

🗳️ 👤 **Dion79** 3 years, 2 months ago

I'd go with A.

<https://www.seagate.com/support/kb/my-system-reported-a-smart-error-on-the-drive-184619en/>

upvoted 4 times

🗳️ 👤 **haazybanj** 3 years, 3 months ago

A is the correct answer

upvoted 3 times

A server administrator has been creating new VMs one by one. The administrator notices the system requirements are very similar, even with different applications. Which of the following would help the administrator accomplish this task in the SHORTEST amount of time and meet the system requirements?

- A. Snapshot
- B. Deduplication
- C. System Restore
- D. Template

Suggested Answer: D

Community vote distribution

D (100%)

🗨️ 👤 **Grumpy_Old_Coot** 8 months, 4 weeks ago

Selected Answer: D

To expand on what Pongsathorn mentioned: If you -clone- a VM, the GUID/etc also gets copied. This will come back to bite you down the road if SCCM or any other package dependent on unique GUID/etc things is used in a management role. Deploying via a Template is effectively the same as installing a brand-new preconfigured VM each time.

upvoted 1 times

🗨️ 👤 **Pongsathorn** 1 year, 8 months ago

Selected Answer: D

TEMPLATE DEPLOYMENT

A clone is an exact image of a system. VM clones are not suited for mass deployment of virtual machines A template is a more generalized image that is intended for multiple uses. Templates act as a baseline image for a VM. While cloned systems can be turned on, templates cannot. They are used for deployment only.

upvoted 3 times

Which of the following steps in the troubleshooting theory should be performed after a solution has been implemented? (Choose two.)

- A. Perform a root cause analysis
- B. Develop a plan of action
- C. Document the findings
- D. Escalate the issue
- E. Scope the issue
- F. Notify the users

Suggested Answer: AC

Community vote distribution

AC (100%)



  **Dion79**  2 years, 10 months ago

The following list represents the basic steps in a troubleshooting methodology:

Identify the problem
 Determine the scope of the problem
 Establish a theory of probable cause/question the obvious
 Test the theory to determine the cause
 Establish a plan of action
 Implement the solution or escalate the issue
 Verify full system functionality
 Implement preventive measures
 Perform a root cause analysis
 Document findings, actions, and outcomes throughout the process.
 upvoted 6 times

  **i_bird** 2 years, 4 months ago

Root cause analysis is part of "implementing preventive measure" cause it helps you in:
 Knowing WHAT happened
 WHY it happened
 and HOW to prevent it from happening again.
 upvoted 2 times

  **Timock** 2 years, 3 months ago



A&C as its obviously state in Dions list.
 upvoted 2 times

  **Dion79**  2 years, 8 months ago

Selected Answer: AC
 Troubleshooting methodology order. according to CompTIA Server + Study Guide SK0-005
 upvoted 5 times

  **ccoli**  7 months, 1 week ago

I want to do RCA during the implementation of solutions not afterward otherwise how do you know you implemented the right fix if you weren't certain of the cause. Due to this C and F seem right to me, document and then tell people they can use it again.
 upvoted 1 times

  **ITken** 1 year, 3 months ago

I also thought the answers were A and C, but from CompTIA:

<https://www.comptia.org/blog/troubleshooting-methodology>

1. Identify the problem

2. Establish a theory of probable cause
3. Test the theory to determine the cause
4. Establish a plan of action to resolve the problem and implement the solution
5. Verify full system functionality, and, if applicable, implement preventive measures
6. Document findings, actions and outcomes

Root cause should be identified in step #3, and you can't do step #5 where you implement preventive measures without knowing root cause.

However, in the real world, cause isn't necessarily root cause. Cause, in most cases, is be a byproduct of root cause. What you implement to fix the problem isn't always the same root issue that caused the problem. I've been in IT for over 30 years and I've seen many-a-cases where root cause is different than the byproduct issue that the user was dealing with. But we have to think like CompTIA and stick to their theory to the T.

upvoted 2 times

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

Selected Answer: AC

- 1 Identify the problem and determine the scope
- 2 Establish a theory of probable cause (question the obvious)
- 3 Test the theory to determine the cause
- 4 Establish a plan of action to resolve the problem (Notify impacted users is here)
- 5 Implement the solution or escalate
- 6 Verify full system functionality and, if applicable, implement preventive measures
- 7 Perform a root cause analysis
- 8 Document findings, actions, and outcomes throughout the process

upvoted 3 times

🗨️ 👤 **PEsty93** 2 years, 8 months ago

Document is not a last step. You document throughout.

upvoted 1 times

🗨️ 👤 **Dion79** 2 years, 6 months ago

You are correct, but final step according to CompTIA is outcome, finds, and conclusion. You need to know your final answer through all the troubleshooting and documentation that you do.

upvoted 1 times

🗨️ 👤 **Dion79** 2 years, 9 months ago

Answer looks suspect. If you follow the troubleshooting methodology A&C are after Implementing the solution no escalation was needed since the technician found the solution then A - perform a root cause analysis with C documenting findings.

upvoted 2 times

A snapshot is a feature that can be used in hypervisors to:

- A. roll back firmware updates.
- B. restore to a previous version.
- C. roll back application drivers.
- D. perform a backup restore.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **hasquaati** 9 months, 2 weeks ago

Selected Answer: B

Snapshots is not a back up strategy
upvoted 2 times

🗳️ 👤 **Alizade** 1 year, 2 months ago

Selected Answer: B

A snapshot is a feature in hypervisors that captures the state of a virtual machine (VM) at a specific point in time. This includes the VM's configuration, memory, and virtual disk data. Snapshots can be used to restore a VM to a previous version, allowing administrators to revert the VM to a known good state in case of problems, such as failed updates, configuration changes, or application issues.
upvoted 2 times

🗳️ 👤 **Obi_Wan_Jacoby** 1 year, 5 months ago

Selected Answer: B

Answer is B. The one of the differences between a storage snapshot and a backup is that the snapshot is stored at the same location as the original data. Therefore, it depends entirely on the reliability of the source. This means that in case of a disaster or damage to the source data, the storage snapshot will be lost or inaccessible. There is no way to restore if the source gets lost. In addition, snapshots on their own do not have the means to check for corruption or restore capabilities.
upvoted 3 times

🗳️ 👤 **Obi_Wan_Jacoby** 1 year, 5 months ago

Hyper-V snapshots are not a backup alternative

The most important thing that you should know about Hyper-V snapshots is that they are not backups and cannot provide the same level of data protection as backup software. VM snapshots are merely a short-term solution for saving the VM state at a particular point in time but it doesn't actually create a copy of the virtual disk. Hyper-V snapshots cannot protect against issues that might affect the host. If the VM gets damaged, created snapshots will be deleted as a result. Thus, the main virtual disk remains a single point of failure in your environment.
upvoted 2 times

A server administrator needs to deploy five VMs, all of which must have the same type of configuration. Which of the following would be the MOST efficient way to perform this task?

- A. Snapshot a VM.
- B. Use a physical host.
- C. Perform a P2V conversion.
- D. Use a VM template.

Correct Answer: D

Deploying a virtual machine from a template creates a virtual machine that is a copy of the template. The new virtual machine has the virtual hardware, installed software, and other properties that are configured for the template.

Reference:

https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vm_admin.doc/GUID-8254CD05-CC06-491D-BA56-A773A32A8130.html

Prerequisites

You must have the following privileges to deploy a virtual machine from a template:

- **Virtual machine.Inventory.Create from existing** on the data center or virtual machine folder.
- **Virtual machine.Configuration.Add new disk** on the data center or virtual machine folder. Required only if you customize the original hardware by adding a new virtual disk.
- **Virtual machine.Provisioning.Deploy template** on the source template.
- **Resource.Assign virtual machine to resource pool** on the destination host, cluster, or resource pool.
- **Datastore.Allocate space** on the destination datastore.
- **Network.Assign network** on the network to which the virtual machine is assigned. Required only if you customize the original hardware by adding a new network card.
- **Virtual machine.Provisioning.Customize** on the template or template folder if you are customizing the guest operating system.
- **Virtual machine.Provisioning.Read customization specifications** on the root vCenter Server if you are customizing the guest operating system.
- **Virtual machine.Inventory.Create new** on the destination folder or data center.
- **vApp.Import**

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

A global organization keeps personnel application servers that are local to each country. However, a security audit shows these application servers are accessible from sites in other countries. Which of the following hardening techniques should the organization use to restrict access to only sites that are in the same country?

- A. Configure a firewall
- B. Close the unneeded ports
- C. Install a HIDS
- D. Disable unneeded services.

Correct Answer: A

Monitors Network Traffic -

Reference:

<https://www.fortinet.com/resources/cyberglossary/benefits-of-firewall>


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

The Chief Information Officer (CIO) of a datacenter is concerned that transmissions from the building can be detected from the outside. Which of the following would resolve this concern? (Choose two.)

- A. RFID
- B. Proximity readers
- C. Signal blocking
- D. Camouflage
- E. Reflective glass
- F. Bollards

Suggested Answer: CE

Community vote distribution



jagoichi Highly Voted 2 years, 9 months ago

CE

reference CompTia official server +

The building's architecture may integrate other forms of security, including reflective glass, structural designs that block wireless signals, deceptive signage to camouflage the data center, and a lack of windows.

The building and property landscaping design are often tied to the physical security aspects of the data center.

upvoted 8 times

shanebrown Most Recent 9 months, 2 weeks ago

Selected Answer: CE

C & E. The idea is to prevent signals from exiting the building. not reading the signal

upvoted 1 times

a792193 1 year, 6 months ago

Selected Answer: CE

C. Signal blocking

E. Reflective glass

upvoted 2 times

hasquaati 1 year, 9 months ago

Selected Answer: CE

C and E

upvoted 2 times

Alizade 2 years, 2 months ago

Selected Answer: CE

C. Signal blocking

E. Reflective glass

upvoted 2 times

Obi_Wan_Jacoby 2 years, 5 months ago

Selected Answer: CE

I am going with C&E. Signal blocking goes without saying. As for E (Reflective glass) it is noted that reflective glass (mirror glass, tinted glass) heavily absorb wifi signals. Having these on the outside walls of the building would not that single down greatly.

upvoted 2 times

Greedy1985 2 years, 5 months ago

I would say CE. It was a similar question on the CompTia Certmaster Practice Exam for Server+. Reflective Glass and Signal Blocking will prevent transmissions from being read outside of the building.

upvoted 2 times

🗨️ 👤 **Pongsathorn** 2 years, 7 months ago

Selected Answer: BC

Signal Blocking

One technique that can be used to prevent data leakage through radio waves is to use signal blocking materials on walls and windows. This can prevent the reception of any wireless transmissions outside the facility. It is useful to know that by using a high-powered antenna (which is illegal) a hacker can be far away from your building and still receive these signals. While many building materials such as metal and concrete will provide limited protection, you can also use special paint on the walls that will block these signals.

Sometimes emissions coming from the servers themselves can disclose sensitive information. This issue can be addressed by placing the server inside enclosures that can block signals. One example is called a Faraday cage, which implements an outer barrier or coating called a Faraday shield.

upvoted 3 times

🗨️ 👤 **Pongsathorn** 2 years, 7 months ago

Reflective Glass

In areas where sensitive operations are being performed or where sensitive discussions or planning may be taking place, you need to prevent prying eyes from seeing what's going on. When there are windows in the area of concern, you should use reflective glass to prevent the viewing of information that may be written on boards or displayed on screens.

Camouflage

Your datacenter should be located in an isolated area if possible and made indistinguishable from other industrial buildings. CPTED calls for using natural landscape elements for camouflage, such as dense trees or even a mountain range. Finally, a datacenter should be enclosed by a wall or metal fence that is difficult to pass through.

upvoted 1 times

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

Let's start with the basics. What KIND of transmissions is the CIO referring to. It matters. A lot.

upvoted 2 times

A server administrator is configuring the IP address on a newly provisioned server in the testing environment. The network VLANs are configured as follows:

VLAN name	VLAN ID	Gateway IP address	Active switchports
Testing	10	192.168.10.1/24	2, 4, 6, 8, 10, 12, 14, 18
Production	20	192.168.20.1/24	3, 5, 7, 9, 11, 13, 15, 17
Administration	30	192.168.30.1/24	1, 24

The administrator configures the IP address for the new server as follows:

IP address: 192.168.1.1/24 -

Default gateway: 192.168.10.1 -


A ping sent to the default gateway is not successful. Which of the following IP address/default gateway combinations should the administrator have used for the new server?

- A. IP address: 192.168.10.2/24 Default gateway: 192.168.10.1
- B. IP address: 192.168.1.2/24 Default gateway: 192.168.10.1
- C. IP address: 192.168.10.3/24 Default gateway: 192.168.20.1
- D. IP address: 192.168.10.24/24 Default gateway: 192.168.30.1

Suggested Answer: A

Community vote distribution


A (100%)

 **Drewid91** 9 months, 1 week ago

Selected Answer: A

Very clearly should be A, They should be on the same subnet.

upvoted 4 times

 **Timock** 9 months, 2 weeks ago


Selected Answer: A

The question states a new server in the Testing LAN. The IP should be on the same LAN. 192.168.10.1/24 is the Gateway IP address from the table. The server cannot use that 10.1 but CAN use 10.2 So A.

VLANs allow network administrators to automatically limit access to a specified group of users by dividing workstations into different isolated LAN segments. When users move their workstations, administrators don't need to reconfigure the network or change VLAN groups.

<https://www.n-able.com/blog/what-are-vlans>

upvoted 4 times

 **jagoichi** 9 months, 2 weeks ago

A

needs to be on same network as gateway

upvoted 2 times

A server administrator is configuring a new server that will hold large amounts of information. The server will need to be accessed by multiple users at the same time. Which of the following server roles will the administrator MOST likely need to install?

- A. Messaging
- B. Application
- C. Print
- D. Database

Correct Answer: D

Few people are expected to use the database at the same time and users don't need to customize the design of the database.

Reference:

<https://support.microsoft.com/en-us/office/ways-to-share-an-access-desktop-database-03822632-da43-4d8f-ba2a-68da245a0446>

To share a database by using a shared folder:

1. In a home or small business environment, share a folder with specific people. For more information, see [File sharing over a network in Windows 10](#).
2. Make sure that Access is set to open in shared mode on all of the users' computers. This is the default setting, but you should check to be sure — if a user opens the database in exclusive mode, it will interfere with data availability.
 - a. Start Access and under **File**, click **Options**.
 - b. In the **Access Options** box, click **Client Settings**.
 - c. In the **Advanced** section, under **Default open mode**, select **Shared**, click **OK**, and then exit Access.
3. Copy the database file to the shared folder. After you copy the file, make sure that the file attributes are set to allow read/write access to the database file. Users must have read/write access to use the database.
4. On each user's computer, create a shortcut to the database file. For more information, see [Create a desktop shortcut for an Office program or file](#).

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

Users at a company work with highly sensitive data. The security department implemented an administrative and technical control to enforce least-privilege access assigned to files. However, the security department has discovered unauthorized data exfiltration. Which of the following is the BEST way to protect the data from leaking?

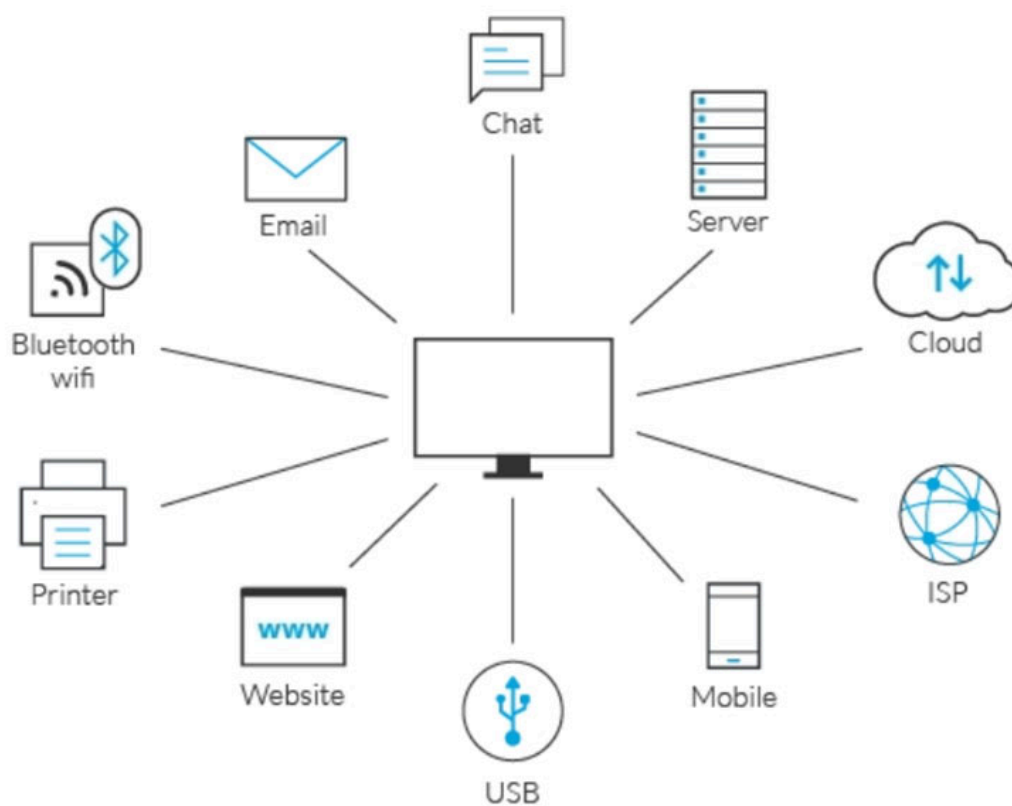
- A. Utilize privacy screens.
- B. Implement disk quotas.
- C. Install a DLP solution.
- D. Enforce the lock-screen feature.

Correct Answer: C

Components of a Data Loss Solution

Reference:

<https://www.imperva.com/learn/data-security/data-loss-prevention-dlp/>



DLP: Preventing data loss from various sources


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

A server administrator needs to create a new folder on a file server that only specific users can access. Which of the following BEST describes how the server administrator can accomplish this task?

- A. Create a group that includes all users and assign it to an ACL.
- B. Assign individual permissions on the folder to each user.
- C. Create a group that includes all users and assign the proper permissions.
- D. Assign ownership on the folder for each user.

Suggested Answer: B

Community vote distribution



🗳️ 👤 **Spacecluster** Highly Voted 🍌 2 years, 8 months ago

The question is confusing, only sensible answer is C assuming "all users" are "specific users"
upvoted 5 times

🗳️ 👤 **PaytoPlay** Most Recent 🔍 5 months, 2 weeks ago

Selected Answer: C

A Group "Share" is created for members to access then NTFS Group permission's are created for granular access to the folders inside the share.
upvoted 1 times

🗳️ 👤 **So_fly58** 5 months, 4 weeks ago

Selected Answer: C

Always modify by group, not by user - whenever possible.
upvoted 1 times

🗳️ 👤 **ompk** 8 months, 2 weeks ago

Assuming the company has worldwide users (more than 2k) and some specific users whether local or abroad, needs to access the file for their project. It is assumed that file is related to one project which should be accessed by the users who are working on the same project but may be belonged to different departments as well as cities or may be countries they living in.
I think C is the right answer where you need to create a group named " The Project" so admin can easily add those users to the group to access the file
upvoted 1 times

🗳️ 👤 **[Removed]** 10 months, 3 weeks ago

Selected Answer: C

It has to be bad wording and they mean "all those specific users", because it would be really weird to tell admins that assigning individual users to a folder is the best option, no one does this, unless there's just one person.
upvoted 1 times

🗳️ 👤 **Fart2023** 11 months ago

Selected Answer: C

C Assuming that this is just a "Bad English" question.
upvoted 1 times

🗳️ 👤 **Fakecon** 1 year ago

Selected Answer: C

Regardless of reason.
The worst is when you have to clear individual users from any folder or file left by previous employee.

Best practice is always to place user in group
upvoted 1 times

🗳️ 👤 **Dingos** 2 years, 1 month ago

Selected Answer: B

I just can not understand CompTIA logic to select this answers so only option is B.
C is not correct because word trap (specific in question, all in C answer)

upvoted 1 times

🗨️ 👤 **agabeen** 2 years, 4 months ago

IT's C..

clear for me

upvoted 1 times

🗨️ 👤 **RSMCT2011** 2 years, 8 months ago

Selected Answer: B

Best answer that meet objective is B:

It is NOT a best practice to assign permission to individual user BUT it is BEST answer among others

upvoted 3 times

🗨️ 👤 **Dingos** 2 years, 1 month ago

Sh*t I just changed my mind, yes I am going for B and I hate even possibility that is something that CompTIA think it is A OK!

upvoted 1 times

🗨️ 👤 **Dingos** 2 years, 1 month ago

Only Sh*t I see here from CompTIA is that they put vague questions and even more vague answers, in question is SPECIFIC users in least wrong answer there is word ALL users.

When I remember I was cursing CISCO for trick questions, now when my MCSA needs replacement I miss normal trick questions.

upvoted 1 times

🗨️ 👤 **Dingos** 2 years, 1 month ago

I disagree. Giving explicit rights is just opening pandora's box of troubles if you have many users on domain. I prefer to make group even there is only one current user on some FS folder. That way you know every access right directly from AD and easy to manage when it comes to position or function change.

So I go with C

upvoted 1 times

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

How are the terms "all users" and "specific users" synonymous. None of the answers shown are correct.

upvoted 3 times

A technician has received multiple reports of issues with a server. The server occasionally has a BSOD, powers off unexpectedly, and has fans that run continuously. Which of the following BEST represents what the technician should investigate during troubleshooting?

- A. Firmware incompatibility
- B. CPU overheating
- C. LED indicators
- D. ESD issues

Suggested Answer: B

Unexpected shutdowns. If the system is randomly shutting down or rebooting, the most likely cause is a heat problem.

Reference:

<https://www.microsoftpressstore.com/articles/article.aspx?p=2224043&seqNum=3>

Common Symptoms

The following are some common symptoms and possible causes related to the CPU or RAM:

- **Unexpected shutdowns.** If the system is randomly shutting down or rebooting, the most likely cause is a heat problem. Check the ventilation and clean out the fans.
- **System lockups.** When a computer stops responding to inputs from the keyboard or mouse, technicians refer to it as frozen or locked up. This can also be due to heat issues. Check the ventilation.
- **Continuous reboots.** In some cases, a hardware issue can prevent the system from booting completely. It starts, gets so far, and then resets itself. This is more common after a faulty software update, but it can be due to a hardware problem. If you've just replaced hardware, double-check your steps. If that isn't the issue, boot into Safe Mode and troubleshoot the operating system using the steps provided in Chapter 17, "Troubleshooting Windows Operating Systems."

Community vote distribution

B (100%)

EngAbood 11 months, 3 weeks ago

Selected Answer: B

correct

upvoted 1 times

Spacecluster 2 years, 2 months ago

Selected Answer: B

Fans run continuously, that means CPU is overheating.

upvoted 1 times

Spacecluster 2 years, 2 months ago

Fans run continuously, that means CPU is overheating.

upvoted 1 times

Which of the following would a systems administrator implement to ensure all web traffic is secure?


- A. SSH
- B. SSL
- C. SMTP
- D. PGP

Correct Answer: B

Secure Sockets Layer (SSL): SSL and its successor Transport Layer Security (TLS) enable client and server computers to establish a secure connection session and manage encryption and decryption activities.

Reference:

<https://paginas.fe.up.pt/~als/mis10e/ch8/chpt8-4bullettext.htm>



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

An administrator is configuring a server to communicate with a new storage array. To do so, the administrator enters the WWPN of the new array in the server's storage configuration. Which of the following technologies is the new connection using?

- A. iSCSI
- B. eSATA
- C. NFS
- D. FCoE

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **agabeen** 10 months, 1 week ago

Selected Answer: D

WWPN mean (world Wide Port Number) which used only of fiber channel in SAN
upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 11 months, 2 weeks ago

Selected Answer: D

Answer is D. Page 142 in the "CompTia Server+ Cerification Exam Guide"
upvoted 2 times

🗳️ 👤 **Pongsathorn** 1 year, 2 months ago

Selected Answer: D

Should be D
Fibre Channel

A very common interface for storage in servers is a Fibre Channel interface. These are used to connect devices in a high-speed fiber storage network. These networks typically use a fiber switch with devices connected to the switch using Fibre Channel interfaces. Servers will require some implementation of a fiber HBA. Each HBA has a unique World Wide Name (WWN), which like a MAC address uses an organizationally unique identifier (OUI) assigned by the IEEE. An example of a fiber switch to which the cables from the HBAs on the servers would attach.

upvoted 3 times

🗳️ 👤 **Timock** 1 year, 3 months ago

A WWPN is a World Wide Port Name; a unique identifier for each Fibre Channel port presented to a Storage Area Network (SAN). Each port on an IBM Storage Device has a unique and persistent WWPN.

https://en.wikipedia.org/wiki/World_Wide_Port_Name#

upvoted 1 times

🗳️ 👤 **jagoichi** 1 year, 3 months ago

Should be D

A World Wide Name (WWN) or World Wide Identifier (WWID) World wide port name is a unique identifier used in storage technologies including Fibre Channel, Parallel ATA, Serial ATA, NVM Express, SCSI and Serial Attached SCSI (SAS).[1]

upvoted 1 times

HOTSPOT -


A systems administrator deployed a new web proxy server onto the network. The proxy server has two interfaces: the first is connected to an internal corporate firewall, and the second is connected to an internet-facing firewall. Many users at the company are reporting they are unable to access the Internet since the new proxy was introduced. Analyze the network diagram and the proxy server's host routing table to resolve the Internet connectivity issues.

INSTRUCTIONS -

Perform the following steps:

1. Click on the proxy server to display its routing table.
2. Modify the appropriate route entries to resolve the Internet connectivity issue.

If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.



Hot Area:

Proxy Server Routing Table

Destination	Netmask	Gateway	Interface
0.0.0.0	0.0.0.0		
		192.168.3.0	192.168.4.1
		192.168.4.0	192.168.1.1
		192.168.1.1	192.168.3.0
		192.168.2.0	192.168.1.0
		192.168.1.0	192.168.2.2
		192.168.4.1	0.0.0.0
		192.168.2.1	192.168.3.1
		0.0.0.0	255.255.255.0
		192.168.3.1	192.168.3.2
		255.255.255.0	192.168.4.0
		192.168.3.2	192.168.2.1
		192.168.2.2	192.168.2.0
192.168.1.0	255.255.255.0		
		192.168.3.0	192.168.4.1
		192.168.4.0	192.168.1.1
		192.168.1.1	192.168.3.0
		192.168.2.0	192.168.1.0
		192.168.1.0	192.168.2.2
		192.168.4.1	0.0.0.0
		192.168.2.1	192.168.3.1
		0.0.0.0	255.255.255.0
		192.168.3.1	192.168.3.2
		255.255.255.0	192.168.4.0
		192.168.3.2	192.168.2.1
		192.168.2.2	192.168.2.0

Suggested Answer:

Proxy Server Routing Table

Destination	Netmask	Gateway	Interface
0.0.0.0	0.0.0.0		
		192.168.3.0	192.168.4.1
		192.168.4.0	192.168.1.1
		192.168.1.1	192.168.3.0
		192.168.2.0	192.168.1.0
		192.168.1.0	192.168.2.2
		192.168.4.1	0.0.0.0
		192.168.2.1	192.168.3.1
		0.0.0.0	255.255.255.0
		192.168.3.1	192.168.3.2
		255.255.255.0	192.168.4.0
		192.168.3.2	192.168.2.1
		192.168.2.2	192.168.2.0
192.168.1.0	255.255.255.0		
		192.168.3.0	192.168.4.1
		192.168.4.0	192.168.1.1
		192.168.1.1	192.168.3.0
		192.168.2.0	192.168.1.0
		192.168.1.0	192.168.2.2
		192.168.4.1	0.0.0.0
		192.168.2.1	192.168.3.1
		0.0.0.0	255.255.255.0
		192.168.3.1	192.168.3.2
		255.255.255.0	192.168.4.0
		192.168.3.2	192.168.2.1
		192.168.2.2	192.168.2.0

 **fluke92** 7 months, 1 week ago

Default Route (0.0.0.0/0):

Gateway: 192.168.3.1

Interface: 192.168.3.2

This route ensures all traffic destined for external networks is correctly routed through the internet-facing firewall, which is crucial for Internet access.

Internal Network Route (192.168.1.0/24):

Gateway: 192.168.2.1

Interface: 192.168.2.2

This route ensures proper communication within the internal network through the internal firewall.

upvoted 1 times

 **Fart2023** 11 months, 3 weeks ago

Default Route:

Destination: 0.0.0.0

Netmask: 0.0.0.0

Gateway: 192.168.3.1

Interface: 192.168.3.2

Destination: 192.168.1.0

Netmask: 255.255.255.0

Gateway: 192.168.2.1



Interface: 192.168.2.2

upvoted 1 times

  **gingasaurusrex** 2 years, 4 months ago

the subnet is /24 meaning anything in the last digits are the only things changed, I believe the answer that was provided is correct

upvoted 4 times

  **jagoichi** 2 years, 9 months ago

Default routes route to internet

0.0.0.0 0.0.0.0 -> 192.168.4.1 via interface 192.168.3.1

192.168.1.0 255.255.255.0 -> 192.168.1.1 via interface 192.168.2.1

upvoted 3 times

  **Destructo** 2 years, 3 months ago

when we are configuring our default gateway we always set it on the interface closest to your network. the given answers are probably correct. you were on the right track but just that the gateway is on the inside not out.

upvoted 1 times

A systems administrator needs to configure a new server and external storage for a new production application environment. Based on end-user specifications, the new solution needs to adhere to the following basic requirements:

1. The OS must be installed in a separate disk partition. In case of hard drive failure, it cannot be affected.
2. Application data IOPS performance is a must.
3. Data availability is a high priority, even in the case of multiple hard drive failures.

Which of the following are the BEST options to comply with the user requirements? (Choose three.)

- A. Install the OS on a RAID 0 array.
- B. Install the OS on a RAID 1 array.
- C. Configure RAID 1 for the application data.
- D. Configure RAID 5 for the application data.
- E. Use SSD hard drives for the data application array.
- F. Use SATA hard drives for the data application array.
- G. Use a single JBOD for OS and application data.

Suggested Answer: BCE

Community vote distribution

BCE (100%)

🗳️ 👤 **Kabbath1986** 2 months, 1 week ago

Selected Answer: BDE

ChatGPT

✓ B. Install the OS on a RAID 1 array

✓ D. Configure RAID 5 for the application data

✓ E. Use SSD hard drives for the data application array
upvoted 1 times

🗳️ 👤 **ccoli** 12 months ago

This is an absurdly stupid question. I assume it is a typo and the RAID5 should actually be RAID6. RAID1 can only support multiple drive failures if it is more than 2 disks at which point it is RAID10.

upvoted 2 times

🗳️ 👤 **Kraken84** 1 year, 8 months ago

Selected Answer: BCE

Given the requirements, the best options to comply with the user's specifications are:

B. Install the OS on a RAID 1 array.

E. Use SSD hard drives for the data application array.

Considering the need for high data availability and the risk of multiple drive failures, RAID 1 (mirroring) with more than two drives or RAID 10 (a combination of RAID 1 and RAID 0) would be more suitable than RAID 5. Since RAID 10 is not an option here and RAID 1 provides redundancy, C. Configure RAID 1 for the application data would be the third best choice.

upvoted 1 times

🗳️ 👤 **Pongsathorn** 2 years, 8 months ago

Selected Answer: BCE

1. The OS must be installed in a separate disk partition. In case of hard drive failure, it cannot be affected. = RAID 1

2. Application data IOPS performance is a must. = SSD

3. Data availability is a high priority, even in the case of multiple hard drive failures. = RAID 1, RAID 6 writes parity information across the drives as is done in RAID 5, but it writes two stripes, which allows the system to recover from two drive failures whereas RAID 5 cannot.



upvoted 1 times

🗳️ 👤 **Spacecluster** 2 years, 8 months ago

Selected Answer: BCE

Correct answer is B C E

upvoted 2 times

  **Spacecluster** 2 years, 8 months ago

RAID 5 doesn't support multiple drive failure



Answer is B C E

upvoted 3 times

  **ccoli** 12 months ago

Neither does RAID1

upvoted 2 times

  **[Removed]** 10 months, 3 weeks ago

RAID1 can absolutely lose multiple disks if you have more than 2, which you can.

Given the pool of answers, BCE are the correct ones.

upvoted 1 times

A server technician installs a new NIC on a server and configures the NIC for IP connectivity. The technician then tests the connection using the ping command.

Given the following partial output of the ping and ipconfig commands:

```
ipconfig /all
```

```
IPv4 address: 192.168.1.5
```

```
Subnet mask: 255.255.255.0
```

```
Default gateway: 192.168.1.1
```

```
ping 192.168.1.1 with 32 bytes of data:
```

```
Request timed out
```

```
Reply from 192.168.1.1: bytes=32 time<1ms TTL=128
```

```
Request timed out
```

```
Reply from 192.168.1.1: bytes=32 time<1ms TTL=128
```

Which of the following caused the issue?

- A. Duplicate IP address
- B. Incorrect default gateway
- C. DHCP misconfiguration
- D. Incorrect routing table

Suggested Answer: A

Community vote distribution

A (75%)

B (25%)

🗳️ 👤 **fluke92** 7 months, 1 week ago

Selected Answer: A

Intermittent ping responses, as shown in the output, are a strong indicator of a duplicate IP address on the network. When two devices have the same IP address, one might respond to pings while the other might cause conflicts, leading to timeouts.

upvoted 1 times

🗳️ 👤 **Fart2023** 11 months ago

Selected Answer: A

A. Duplicate IP address, there is nothing wrong with the default gateway.

upvoted 1 times

🗳️ 👤 **ccoli** 1 year, 1 month ago

A is correct. With the wrong gateway or even no gateway you could ping 192.168.1.1 assuming something was there it would respond. With a duplicate 1.5 ip on the network you'd only get out half the time because you're competing with it.

upvoted 1 times

🗳️ 👤 **RBL23168** 1 year, 6 months ago

Selected Answer: B

B. Incorrect default gateway

The given output indicates a problem with communication to the default gateway (192.168.1.1), as indicated by the "Request timed out" messages. The partial ping results show that the server is not able to successfully reach the default gateway. This means that the server is unable to send packets to the specified default gateway (192.168.1.1), which could be due to a misconfiguration or a connectivity issue with the gateway itself.

upvoted 1 times

🗳️ 👤 **Kraken84** 1 year, 8 months ago

Selected Answer: A

Wouldn't this essentially be pinging your back door from your front 🚪. Or Vice-a-Versa



upvoted 2 times

🗳️ 👤 **Obi_Wan_Jacoby** 2 years, 5 months ago

Going with answer A. Each time I have ever encountered an IP conflict (two devices having same IP) I have got the one received one lost pings.
upvoted 1 times

  **lordguck** 2 years, 7 months ago

The reasoning seems to be, that the reply packet was routed to the duplicate IP owner and thus not to the original sender.
upvoted 1 times

  **Spacecluster** 2 years, 8 months ago

Don't understand how the answer is A?
Confusing question
upvoted 2 times

A server administrator is swapping out the GPU card inside a server. Which of the following actions should the administrator take FIRST?

- A. Inspect the GPU that is being installed.
- B. Ensure the GPU meets HCL guidelines.
- C. Shut down the server.
- D. Disconnect the power from the rack.

Suggested Answer: B

Community vote distribution

B (82%)

C (18%)

🗳️ 👤 **fluke92** 7 months, 1 week ago

Selected Answer: C

Before performing any hardware changes, including swapping out a GPU card, the first and most critical step is to shut down the server to prevent damage to the hardware and ensure safety. This protects the server's components from electrical surges and avoids potential data loss or corruption.
upvoted 1 times

🗳️ 👤 **error77** 1 year, 4 months ago

Selected Answer: C

If technician is swapping exactly the same card in stock, there's no need to check HCL. C for sure.
upvoted 1 times

🗳️ 👤 **ccoli** 12 months ago

it doesn't specify if the card is the same. Poorly written on purpose to try to collect more fees for test for a cert that is ultimately worthless and not respected in general by anyone who know anything about IT.
upvoted 1 times

🗳️ 👤 **ccoli** 11 months, 1 week ago

I guess since it says he is swapping and not that he is preparing to swap it we're suppose too assume he already checked HCL
upvoted 1 times

🗳️ 👤 **Dingos** 2 years, 1 month ago

one more idiot question that have 3 possible correct answers depending on view.
A - I would before install visually inspect any device that goes inside PC or Server
B- Before going to actual plan to install new GPU I would check HCL compability
C- Of course I would shut down server before opening it or removing from rack

and cherry on cake even D is possible as I most probably had to disconnect power to move server from the rack.....

upvoted 1 times

🗳️ 👤 **Dingos** 2 years, 1 month ago

D option is only one that I am sure it is wrong as I would shut down server before pulling out power. And just because that D is making mention C I will choose C as my answer.
upvoted 1 times

🗳️ 👤 **comptiaboy** 2 years, 2 months ago

Selected Answer: B

Always check HCL before swapping out any component
upvoted 3 times

🗳️ 👤 **gingasaurusrex** 2 years, 2 months ago

Selected Answer: B

it is definitely B, it does not say that they are swapping out the same video card, they said they are swapping a video card for another one, this could be a newer one that might not be compatible. Make sure it is before shutting down the server and carrying out the task is what you should do FIRST
upvoted 3 times

🗳️ 👤 **agabeen** 2 years, 4 months ago

Selected Answer: C

it's mention swapping out inside the server .. means cards already installed but need to swap them ,, so no need to check for HCL ,, need to shut down the server

upvoted 1 times

  **Spacecluster** 2 years, 8 months ago

Selected Answer: B

Hardware must be compatible with HCL (Hardware Compatibility List)

upvoted 3 times

A server administrator must respond to tickets within a certain amount of time. The server administrator needs to adhere to the:

- A. BIA.
- B. RTO.
- C. MTTR.
- D. SLA.

Correct Answer: *D*

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

Which of the following relates to how much data loss a company agrees to tolerate in the event of a disaster?

- A. RTO
- B. MTBF
- C. PRO
- D. MTTR

Suggested Answer: C

Community vote distribution

C (100%)

  **nixonbii** Highly Voted 8 months, 2 weeks ago

There are no correct answers. RPO is the amount of acceptable data loss in the event of a disaster. Bad answers or a typo.
upvoted 7 times

  **Pongsathorn** Highly Voted 7 months, 2 weeks ago

Selected Answer: C

C should read RPO not PRO as RTO is NOT correct.


Recovery point objective (RPO) is the maximum acceptable amount of data loss after an unplanned data-loss incident, expressed as an amount of time.

upvoted 5 times

  **Pongsathorn** Most Recent 8 months ago

Recovery point objective (RPO) is the maximum acceptable amount of data loss after an unplanned data-loss incident, expressed as an amount of time.

upvoted 1 times

  **Timock** 9 months, 2 weeks ago

C should read RPO not PRO as RTO is NOT correct.

Below are some of the factors that can affect RPOs:

Maximum tolerable data loss for the specific organization

Industry-specific factors – businesses dealing with sensitive information such as financial transactions or health records must update more often

Data storage options, such as physical files versus cloud storage, can affect the speed of recovery

The cost of data loss and lost operations

Compliance schemes include provisions for disaster recovery, data loss, and data availability that may affect businesses

The cost of implementing disaster recovery solutions

<https://www.druva.com/blog/understanding-rpo-and-rto/>

upvoted 1 times

  **hakumai** 9 months, 3 weeks ago

Selected Answer: C

Correct is C:RPO


upvoted 4 times

A server administrator is testing a disaster recovery plan. The test involves creating a downtime scenario and taking the necessary steps. Which of the following testing methods is the administrator MOST likely performing?

- A. Backup recovery
- B. Simulated
- C. Tabletop
- D. Live failover

Suggested Answer: B

Community vote distribution



szl0144 Highly Voted 2 years, 3 months ago

Selected Answer: B

I think B seems correct
upvoted 7 times

saltz Most Recent 1 year, 1 month ago

Selected Answer: B

bbbbbbbbbbbbbbbb
upvoted 1 times

Kraken84 1 year, 2 months ago

If the test involves creating a downtime scenario and taking the necessary steps, then the administrator is performing a real-world test, as opposed to a hypothetical or discussion-based test.

The correct answer is:

D. Live failover.

Explanation:

Live failover involves actively switching operations to a disaster recovery site to see if the process works in real-world conditions. It's the most comprehensive test but can be disruptive.

Backup recovery is simply the process of restoring data from backups to ensure that the backup system works.

Simulated testing might involve some technical steps but doesn't typically involve actual downtime or live failover.

Tabletop is a discussion-based scenario where team members walk through the plan and discuss actions in a hypothetical disaster situation without actual technical steps being taken.

upvoted 1 times

error77 10 months, 2 weeks ago

Admin is TESTING the new plan, you don't use live servers for testing purpose, D is wrong, B is correct.
upvoted 1 times

weat 1 year, 8 months ago

From Chat GPT:

A live failover in a disaster recovery plan involves the actual activation of a secondary system or site that takes over the operations and functions of the primary system that has failed. This can involve redirecting traffic, rerouting data, and ensuring business continuity.

A simulation in a disaster recovery plan involves testing the various systems, processes, and procedures involved in the disaster recovery plan without actually activating the disaster recovery plan itself. This can involve creating scenarios and testing how different personnel and systems respond to them.

A tabletop exercise in a disaster recovery plan involves a discussion-based simulation where team members review and discuss the disaster recovery plan in a classroom-style setting. This can be used to identify gaps and potential risks and make improvements to the plan.

Overall, live failover involves actual implementation of disaster recovery plan, simulation involves testing without actual implementation of disaster recovery plan and tabletop exercise involves a discussion-based simulation that helps to identify potential risks and improve the plan.

upvoted 2 times

🗨️ 👤 **Obi_Wan_Jacoby** 1 year, 11 months ago

Selected Answer: C

I believe C is correct. The "Simulation" (answer B) aka "Scenario" would occur within the tabletop exercise. I hate this question "<https://gbq.com/the-benefits-of-conducting-tabletop-testing/>"

upvoted 1 times

🗨️ 👤 **RSMCT2011** 2 years, 2 months ago

Selected Answer: C

In a tabletop test, participants walk through plan activities step by step to demonstrate whether DR team members know their duties in an emergency. A simulation test uses resources such as recovery sites and backup systems in what is essentially a full-scale test without an actual failover.

<https://www.techtarget.com/searchdisasterrecovery/definition/disaster-recovery-plan>

upvoted 1 times

🗨️ 👤 **Timock** 2 years, 3 months ago

A tabletop discussion is a discussion with various heads of depts to create a DRP. A simulated tabletop would be going through step by step that which has already been created to test.

upvoted 1 times

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

Tabletop/Simulated failover—the disaster recovery procedures are implemented on a limited scale. Participants engage in role-playing to ensure comprehension and realism

upvoted 1 times

A technician wants to limit disk usage on a server. Which of the following should the technician implement?

- A. Formatting
- B. Compression
- C. Disk quotas
- D. Partitioning

Correct Answer: C

Reference:

<https://www.digitalcitizen.life/simple-questions-what-are-disk-quotas-how-set-them-windows/>

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

A systems administrator has noticed performance degradation on a company file server, and one of the disks on it has a solid amber light. The administrator logs on to the disk utility and sees the array is rebuilding. Which of the following should the administrator do NEXT once the rebuild is finished?

- A. Restore the server from a snapshot.
- B. Restore the server from backup.
- C. Swap the drive and initialize the disk.
- D. Swap the drive and initialize the array.

Suggested Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **H_A_79** 10 months ago

D is my answer because during initializing the disk it will be formatted
upvoted 1 times

🗨️ 👤 **Alizade** 1 year, 2 months ago

Selected Answer: C

C. Swap the drive and initialize the disk.

Once the rebuild is finished, the systems administrator should swap the problematic drive (the one with the solid amber light) with a new one and initialize the disk. This will ensure that the new disk is properly integrated into the array and will help maintain the health and performance of the company file server.

upvoted 1 times

A server administrator needs to configure a server on a network that will have no more than 30 available IP addresses. Which of the following subnet addresses will be the MOST efficient for this network?

- A. 255.255.255.0
- B. 255.255.255.128
- C. 255.255.255.224
- D. 255.255.255.252

Suggested Answer: C

Community vote distribution

C (100%)

🗉 👤 **Alizade** 8 months, 3 weeks ago

Selected Answer: C

Option C: 255.255.255.224 uses 27 bits for the subnet mask, which provides $2^{(32-27)} - 2 = 30$ available IP addresses. This subnet mask is the most efficient choice for this scenario because it uses the exact number of bits needed to provide enough IP addresses for the network, without wasting any addresses.

upvoted 1 times

A remote physical server is unable to communicate to the network through the available NICs, which were misconfigured. However, the server administrator is still able to configure the server remotely. Which of the following connection types is the server administrator using to access the server?

- A. Out-of-band management
- B. Crash cart access
- C. Virtual administrator console
- D. Local KVM setup
- E. RDP connection

Suggested Answer: A

Community vote distribution

A (100%)

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

Selected Answer: A

not internet connection remotely is Out of band, No internet connection locally is a crash box
upvoted 1 times

🗨️ 👤 **Spacecluster** 1 year, 2 months ago

Selected Answer: A

Correct answer
upvoted 1 times

SIMULATION -

A recent power outage caused email services to go down. A server administrator also received averts from the datacenter's UPS. After some investigation, the server administrator learned that each PDU was rated at a maximum of 12A.

Instructions -

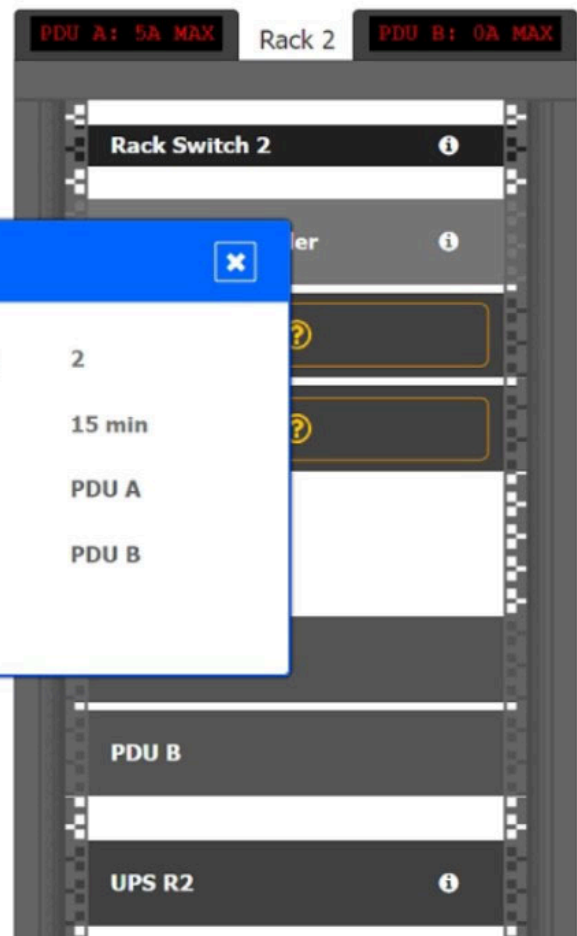
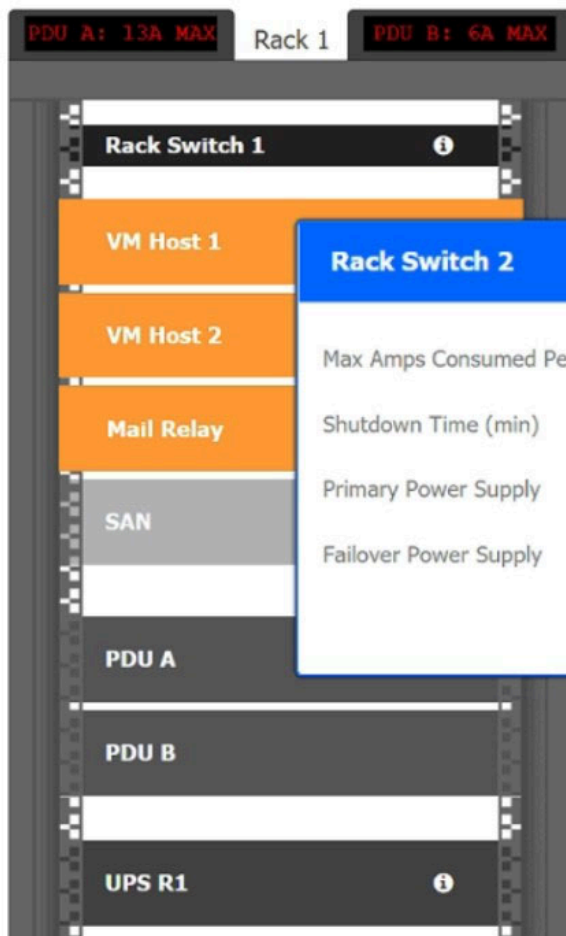
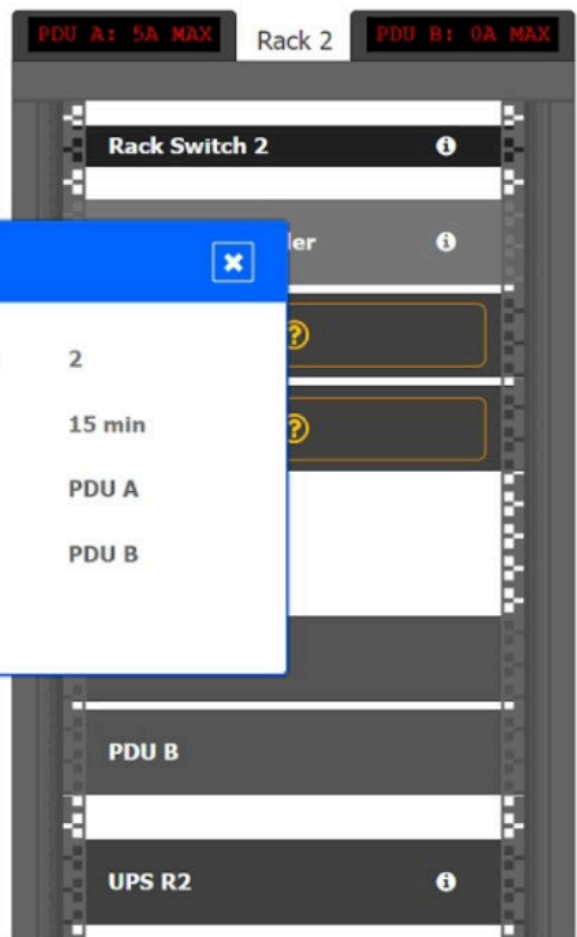
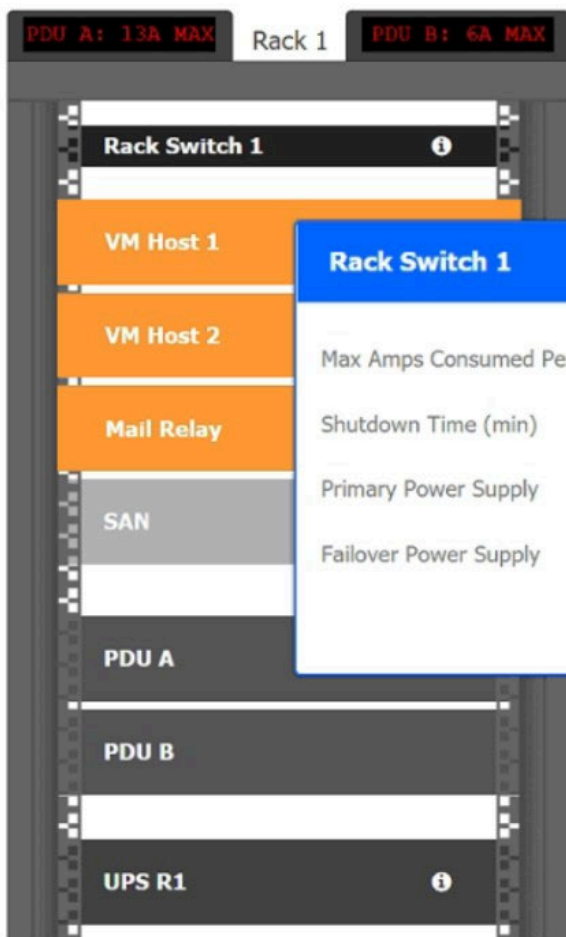
Ensure power redundancy is implemented throughout each rack and UPS alarms are resolved. Ensure the maximum potential PDU consumption does not exceed

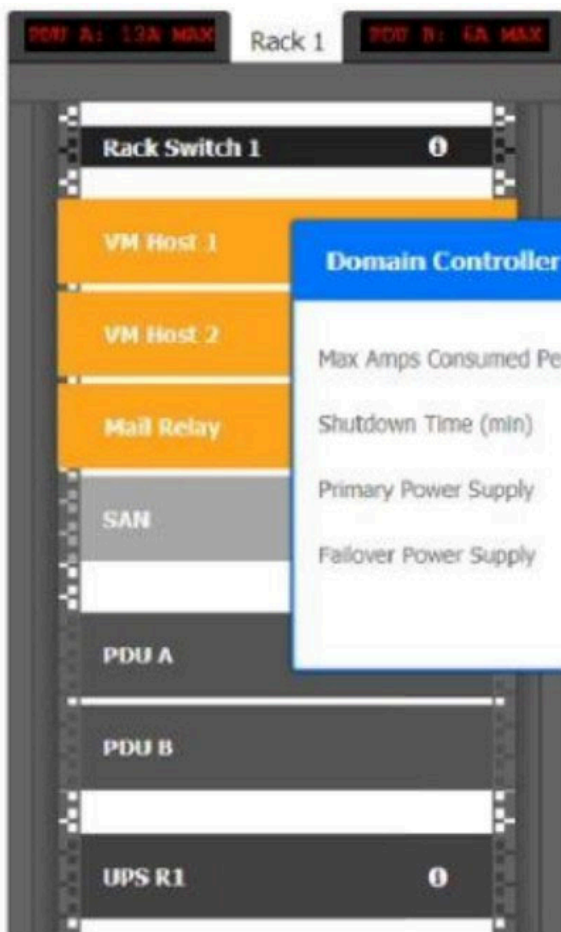
80% (or 9.6A).

- a. PDU selections must be changed using the pencil icon.
- b. VM Hosts 1 and 2 and Mail Relay can be moved between racks. c. Certain devices contain additional details.

If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

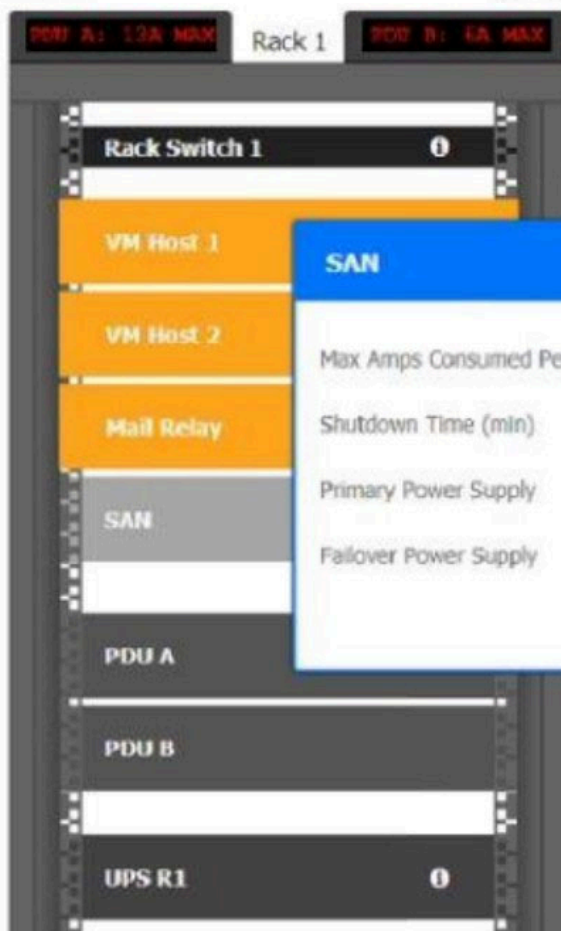






Domain Controller [X]

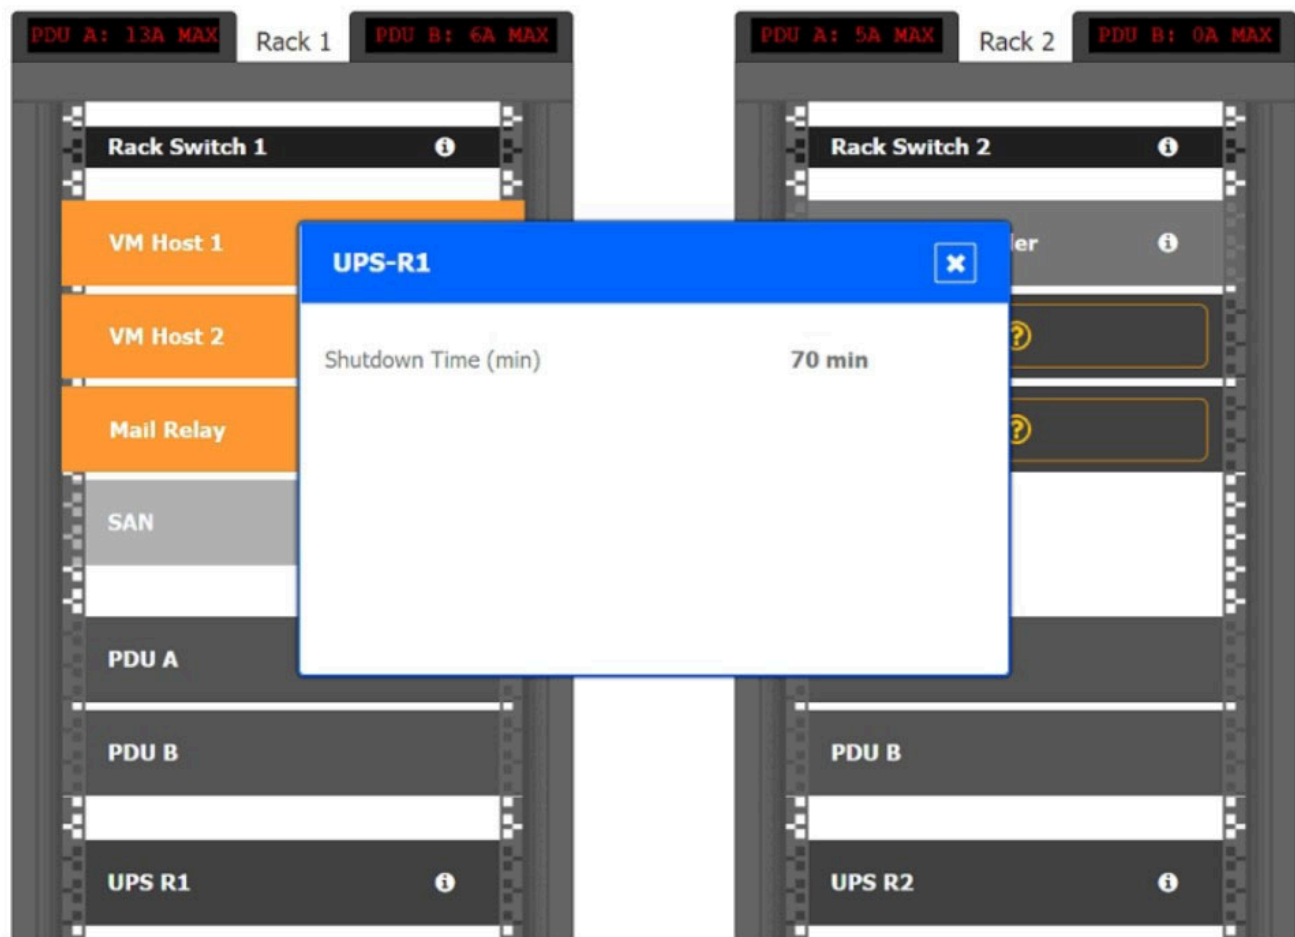
Max Amps Consumed Per Power Supply	3
Shutdown Time (min)	20 min
Primary Power Supply	PDU A
Fallover Power Supply	PDU B

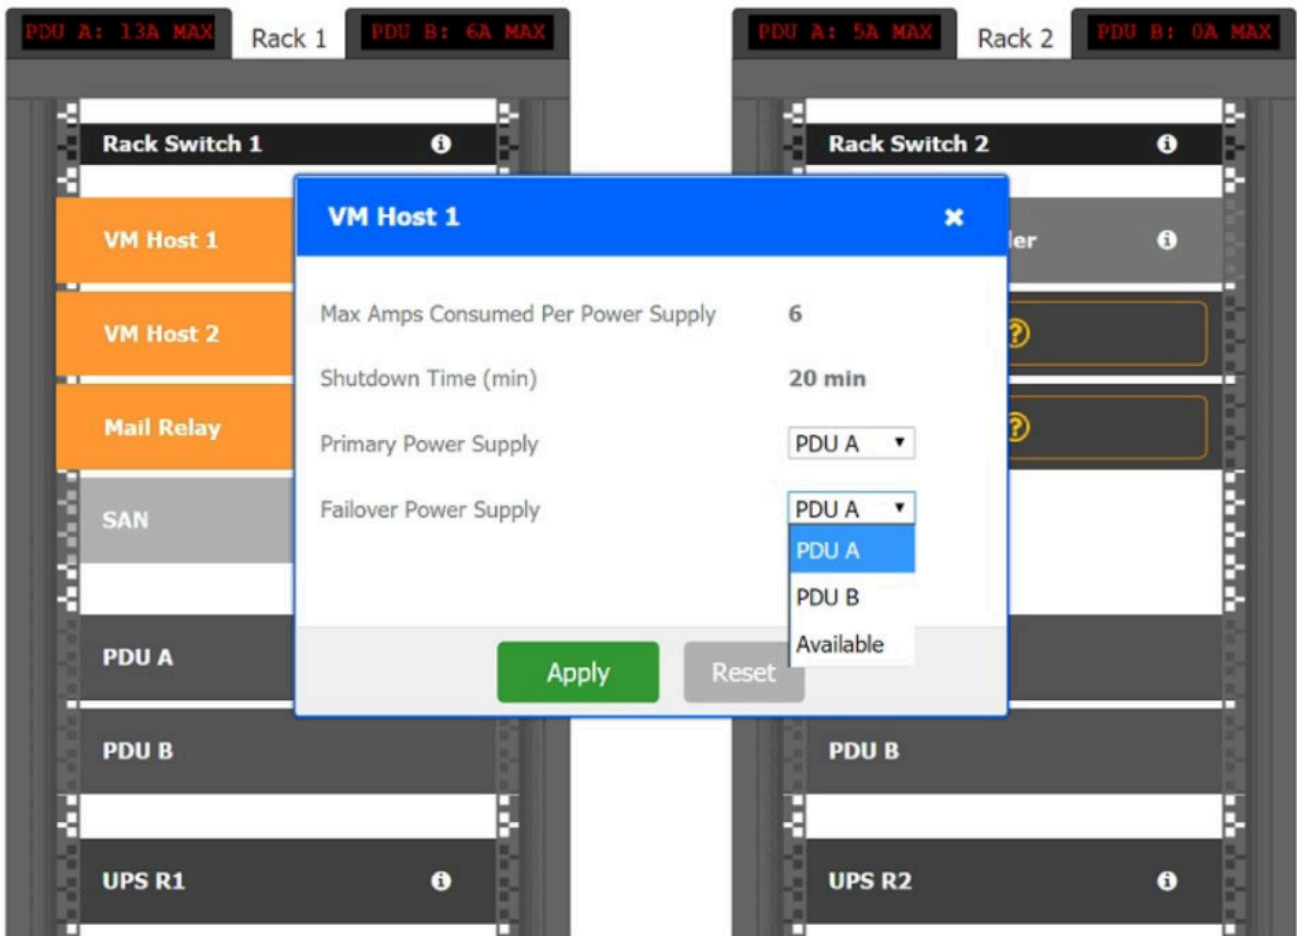
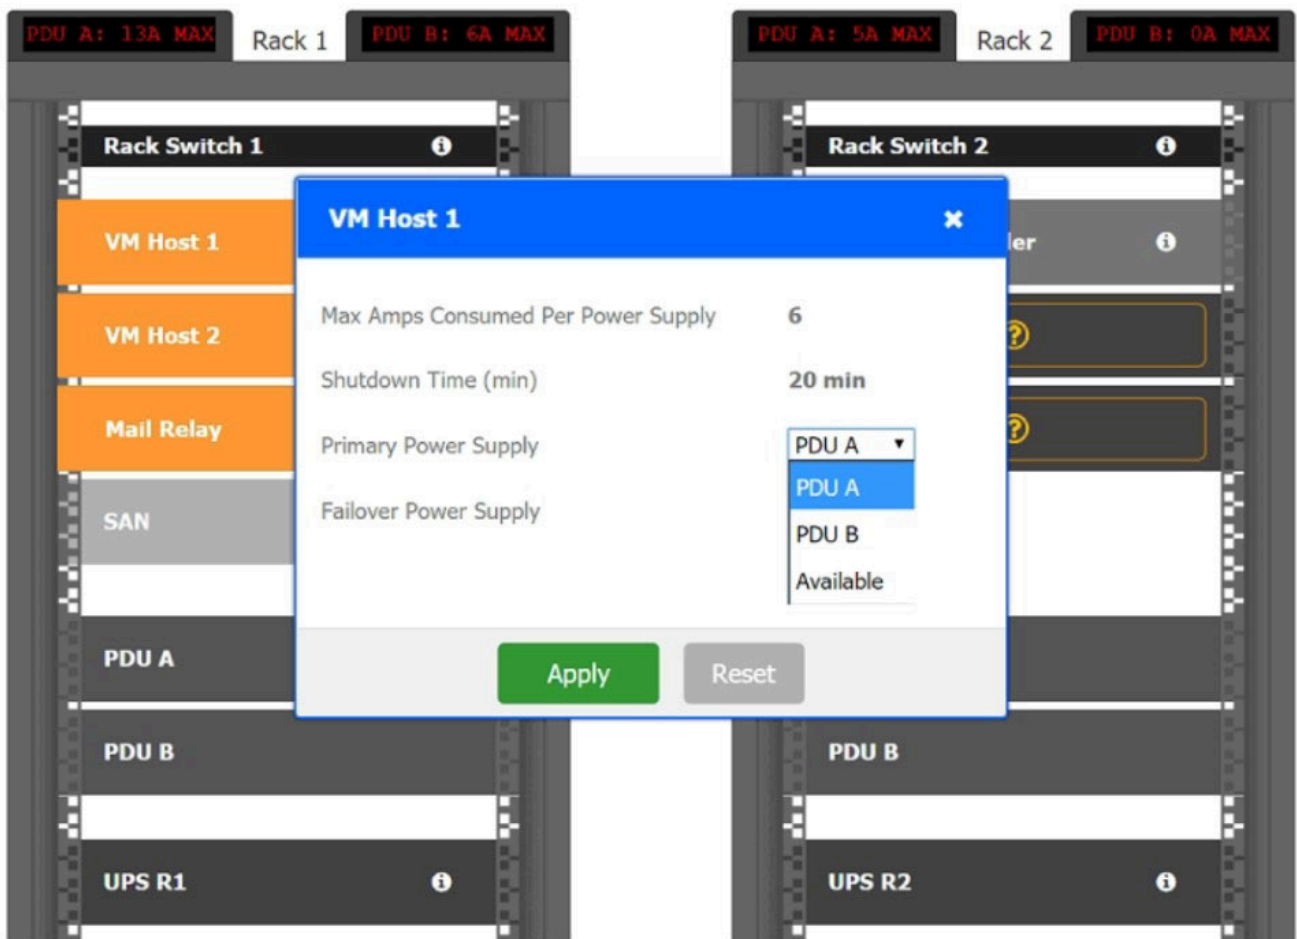


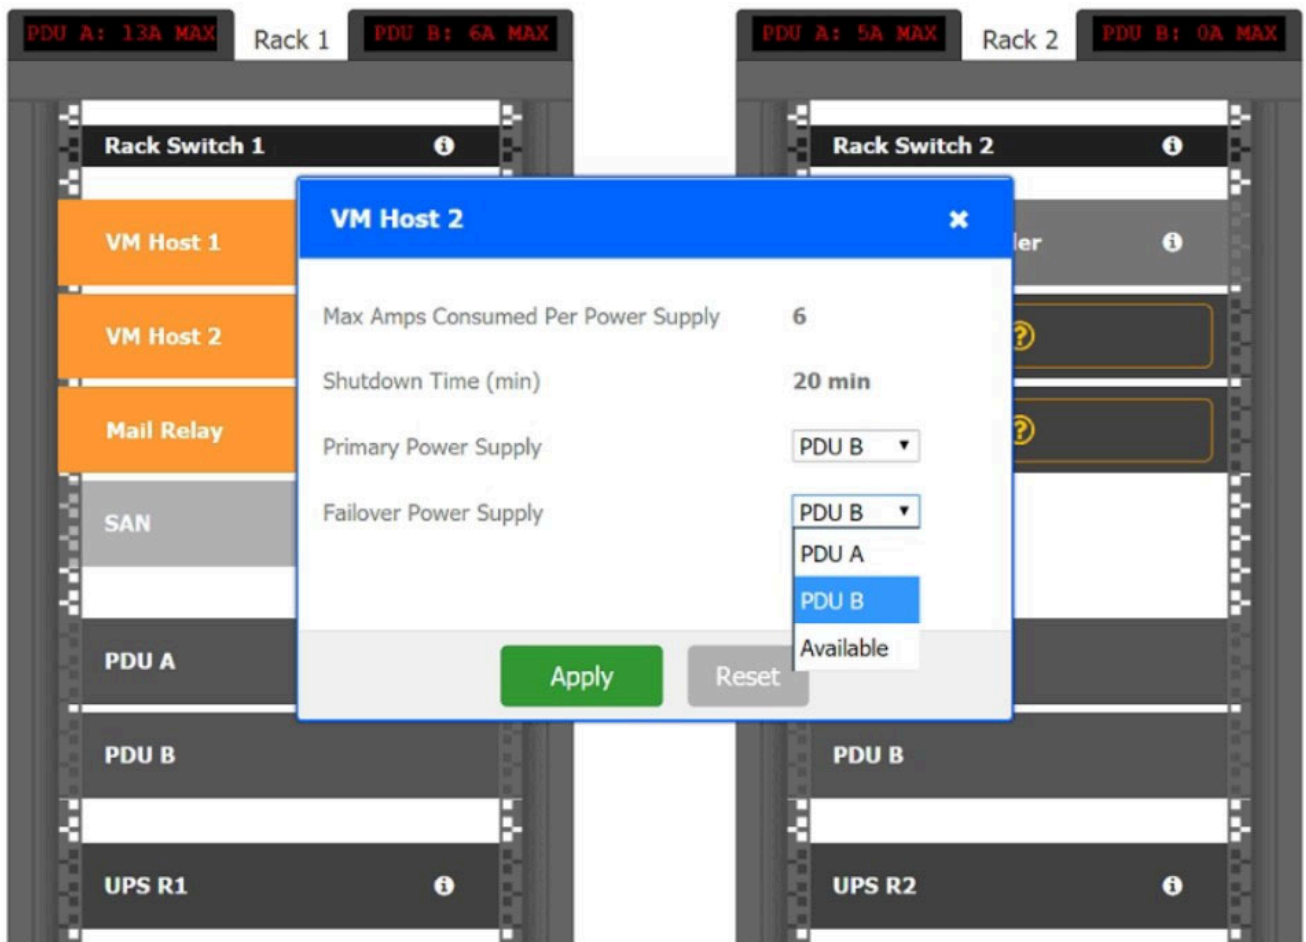
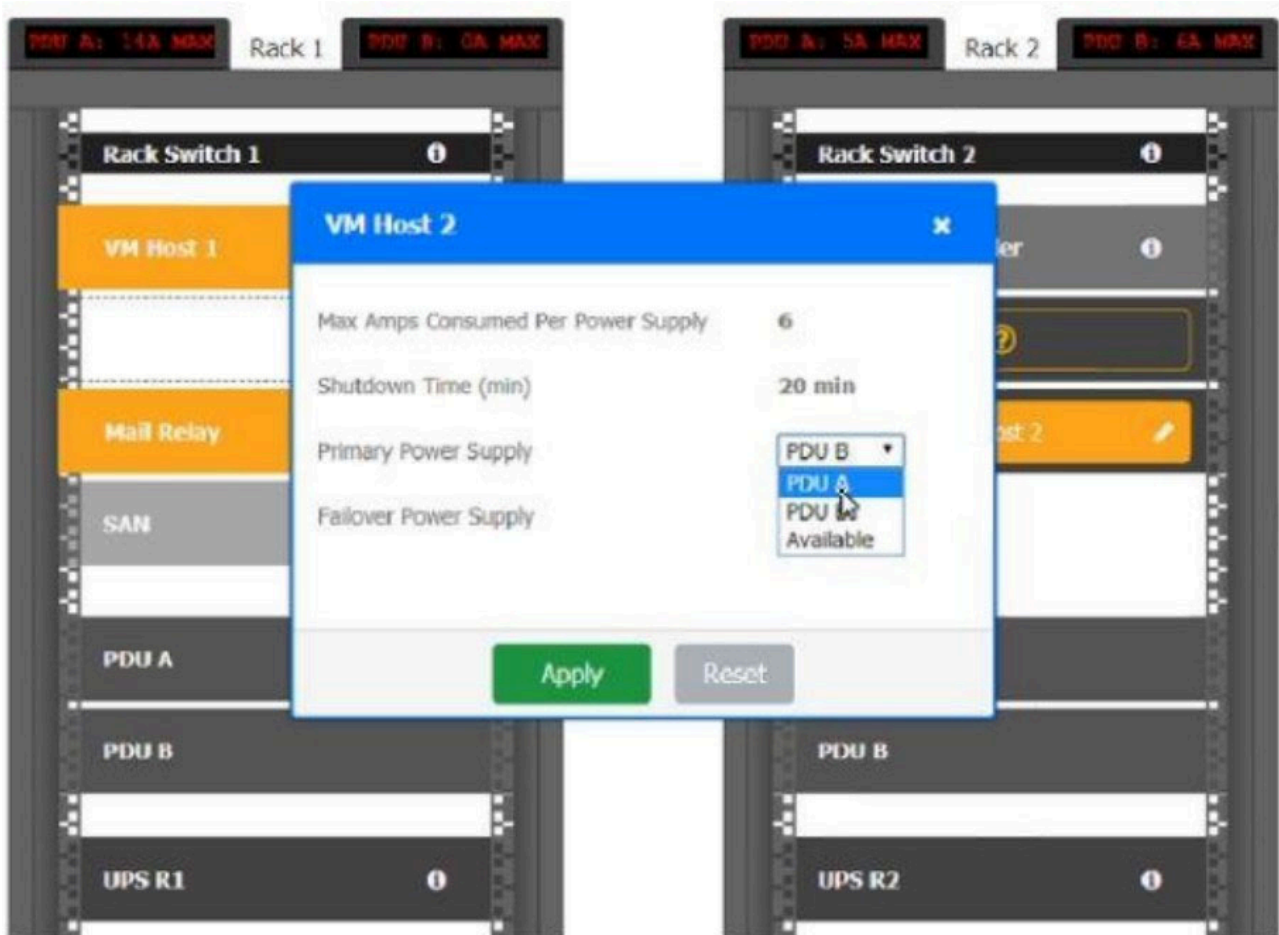
SAN [X]

Max Amps Consumed Per Power Supply	5
Shutdown Time (min)	20 min
Primary Power Supply	PDU A
Fallover Power Supply	PDU B











Suggested Answer: See explanation below.

1. Mailrelay - Keep the mail relay on the existing Rack1 and change the primary power supply to PDU-B and Failover powersupply as PDU B.
 2. Move the VM Host 1 and VM Host2 to Rack2.
- Assign primary power supply PDU A to VM host1
Assign Failover power supply PDU B to VM host1

Assign primary power supply PDU B to VM host2
Assign Failover power supply PDU A to VM host2

🗨️ 👤 **lordguck** Highly Voted 2 years, 7 months ago

1. The proposed solution is wrong, as it seems, you can't change primary/secondary PDU selection on the static installments of the switches, SAN and DC. That gives PDU1 of Rack 2 a remaining capacity of 4.6A, which does not allow HOST1 (6A) to use it as primary power source (9.6A target).
2. There is no solution in case of a PDU failure in Rack 1. The combined power needed will exceed the max 12A in every case.

My solution: VM Host1 -> Rack1, Primary PDU B, secondary A

VM Host2 Primary PDU B, Secondary A

MailRelay -> Rack2 PDU A, Secondary B

This makes

Rack1 13A: PDU1 7A, PDU2 6A <!FAIL in case PDU2 or PDU1 go down.

Rack2 12A: PDU1 6A, PDU2 6A

upvoted 12 times

🗨️ 👤 **Booboo112** 2 years, 4 months ago

So vm host2 goes to rack 2?

upvoted 2 times

🗨️ 👤 **LUCIF4RSD** 8 months, 1 week ago

This is the correct answer. I took the exam yesterday.

Rack1

Vm1 - primary A secondary B

Rack 2

VM 2- primary A secondary B

Mail - both A B can use but (primary B secondary A looks smoother)

upvoted 2 times

🗨️ 👤 **samiraninside** 1 year, 11 months ago

We need to make utilization at 80% that is 9.6 A which seems impossible as static loads are using 12 A from PDU1 itself. PDU2 is still doable. But I am not getting for PDU1

upvoted 1 times

🗨️ 👤 **Fart2023** Most Recent 11 months ago

My answer:

Rack 1:

SW1 (2A), pduA/pduB

VM1 (6A), pduA/pduB

MAIL (1A), pduB/pduA

Rack1 PDU A: 13A MAX, normal load: 8A

Rack1 PDU B: 13A MAX, normal load: 7A

Max shutdown time: 70mins

Rack 2:

SW2 (2A), pduA/pduB

SAN (5A), pduA/pduB

DC (3A), pduB/pduA



VM2 (6A), pduB/pduA

Rack2 PDU A: 12A MAX, normal load: 7A

Rack2 PDU B: 12A MAX, normal load: 9A

Max shutdown time: 65mins

upvoted 2 times

  **error77** 1 year, 4 months ago

Rack1:

SW1 (2A), pduA/pduB

SAN (5A), pduA/pduB

VM1 (6A), pduB/pduA

Rack1 PDU A: 13A MAX, normal load: 7A

Rack1 PDU B: 13A MAX, normal load: 6A

max shutdown time: 55mins

Rack2:

SW2 (2A), pduA/pduB

DC (3A), pduA/pduB

MAIL (1A), pduA/pduB

VM2 (6A), pduB/pduA

Rack2 PDU A: 12A MAX, normal load: 6A

Rack2 PDU B: 12A MAX, normal load: 6A

max shutdown time: 65mins

upvoted 3 times

  **MrS** 1 year, 12 months ago

Click on the panel icon of PDU 1 in Rack 1. Change the PDU selection from A to B. This will balance the power load between UPS A and UPS B and resolve the alarm on UPS A.

Click on the panel icon of PDU 2 in Rack 2. Change the PDU selection from A to B. This will balance the power load between UPS A and UPS B and resolve the alarm on UPS B.

Click on VM Host 1 in Rack 1 and drag it to Rack 2. Drop it below VM Host 2. This will reduce the power consumption of PDU 1 in Rack 1 from 10A to 6A, which is below the 80% threshold of 9.6A.

Click on Mail Relay in Rack 2 and drag it to Rack 1. Drop it below Firewall. This will increase the power consumption of PDU 1 in Rack 1 from 6A to 8A, which is still below the 80% threshold of 9.6A, and reduce the power consumption of PDU 2 in Rack 2 from 10A to 8A, which is also below the 80% threshold of 9.6A.

upvoted 1 times

  **Dingos** 2 years, 1 month ago

What about?

Rack 1: SW1 (15min), VM Host 1 (20 min), Mail Relay (10min), SAN (20min) = 65 min shutdown

Rack 2: SW2 (15min), VM Host 2 (20 min), DC (20min) = 55 min shutdown

Rack 1:

PDU A primary: SW1 (2A), SAN (5A) = 7A



PDU B primary: Mail Relay (1A), VM Host 1 (6A) = 7A

Rack 2:

PDU A primary: SW2 (2A), DC (3A) = 5A

PDU B primary: VM Host 2 (6A) = 6A

upvoted 2 times

  **Dingos** 2 years, 1 month ago

Rack 1:

PDU A primary: SAN (5A), SW1 (2A)

PDU A secondary: Mail Relay (1A), VM Host 1 (6A)

PDU B primary: SAN (5A), SW1 (2A)

PDU B secondary: Mail Relay (1A), VM Host 1 (6A)

Rack 2:



PDU A primary: DC (3A), SW2 (2A)

PDU A secondary: VM Host 2 (6A)

PDU B primary: DC (3A), SW2 (2A)

PDU B secondary: VM Host 2 (6A)

upvoted 1 times

  **error77** 1 year, 4 months ago

Your PDU in rack 1 would exceed 13A (14A), if one of the PDUs fail. Move the mail server to rack 2 to balance the load.

upvoted 1 times

  **billysunshine** 2 years, 1 month ago

what answer ARE comptia looking for in this sim!?

upvoted 2 times

A server administrator is creating a new server that will be used to house customer sales records. Which of the following roles will MOST likely be installed on the server?

- A. Print
- B. File
- C. Database
- D. Messaging

Correct Answer: C

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

A server administrator is deploying a UPS at a datacenter and notices the device is making an audible beep every few seconds. Which of the following is the MOST likely cause?

- A. The battery cable is disconnected
- B. The device is plugged into the wrong outlet
- C. The battery is faulty
- D. The serial cable is not connected

Suggested Answer: C

Community vote distribution

A (100%)

🗳️ 👤 **So_fly58** 5 months, 4 weeks ago

Selected Answer: C

C: Faulty battery.

just replaced one of these this week.

upvoted 1 times

🗳️ 👤 **fluke92** 7 months, 1 week ago

Selected Answer: A

An audible beep from a UPS every few seconds usually indicates a problem with the battery connection or power source. The most common cause is that the battery cable is disconnected, meaning the UPS cannot access its battery backup. This issue often occurs after installation or maintenance when the battery is not properly reconnected.

upvoted 1 times

🗳️ 👤 **kx7tg4xu** 10 months ago

Selected Answer: A

I think A is the answer.

UPSs make periodic sounds, usually to indicate some problem or anomaly.

One of the most common causes is a battery-related problem.

In this situation, it is wise to first check for the simplest and most likely cause, a poor battery cable connection, since the newly installed UPS is sounding the warning tone. This is an easy fix and often solves the problem.

upvoted 1 times

🗳️ 👤 **Timock** 2 years, 9 months ago

If you've lost power, it's beeping to let you know that the battery is in use, and that you should save your work and shut down your computer. A constant beep (every second or two, and never stopping) generally means the UPS is very low on battery power, and you should shut down immediately.

Of the following options it is the most likely although the beeps really mean there is a power issue but from which direction... unplugged? bad battery? low battery?

<https://www.wingswept.com/uninterruptible-power-supply-beeping/#>

upvoted 2 times

A server technician is installing a new server OS on legacy server hardware. Which of the following should the technician do FIRST to ensure the OS will work as intended?

- A. Consult the HCL to ensure everything is supported
- B. Migrate the physical server to a virtual server
- C. Low-level format the hard drives to ensure there is no old data remaining
- D. Make sure the case and the fans are free from dust to ensure proper cooling

Suggested Answer: A

Community vote distribution

A (100%)

 **Timock** Highly Voted 9 months, 2 weeks ago

Selected Answer: A

Since this is legacy server hardware it would be important to make sure that it can handle/support a particular O/S before attempting an install. The rest of the answers have little to nothing to do with an install an O/S on legacy hardware.

upvoted 5 times

An administrator is researching the upcoming licensing software requirements for an application that usually requires very little technical support. Which of the following licensing models would be the LOWEST cost solution?

- A. Open-source
- B. Per CPU socket
- C. Per CPU core
- D. Enterprise agreement

Suggested Answer: A

  **Dlam** 6 months, 3 weeks ago

Open source licenses are licenses that comply with the Open Source Definition – in brief, they allow software to be freely used, modified, and shared. To be approved by the Open Source Initiative (also known as the OSI) a license must go through the Open Source Initiative's license review process.
upvoted 1 times

Which of the following commands should a systems administrator use to create a batch script to map multiple shares?

- A. nbtstat
- B. net use
- C. tracert
- D. netstat

Suggested Answer: B

Community vote distribution

B (100%)

 **gingasaurusrex** 8 months, 2 weeks ago

Selected Answer: B

The command that a systems administrator should use to create a batch script to map multiple shares is option B: net use.

The "net use" command is used to connect or disconnect a computer from a shared resource, such as a network drive or a printer. It can be used to map multiple shares in a batch script, allowing the administrator to automate the process of connecting to multiple network resources.

Option A (nbtstat) is a command-line tool that displays information about the NetBIOS over TCP/IP (NetBT) protocol, such as the NetBIOS name table, the NetBIOS name cache, and the NetBIOS scope ID.

Option C (tracert) is a command-line tool that traces the route that packets take from one network node to another, showing the IP addresses of the routers and the time it takes for the packets to travel to each router.

Option D (netstat) is a command-line tool that displays active TCP connections, listening ports, and network statistics such as the number of packets sent and received, the number of errors, and the number of connections in various states.

upvoted 3 times

An administrator is troubleshooting a failure in the datacenter in which a server shut down/powered off when utility power was lost. The server had redundant power supplies. Which of the following is the MOST likely cause of this failure?

- A. The UPS batteries were overcharged
- B. Redundant power supplies require 220V power
- C. Both power supplies were connected to the same power feed
- D. The power supplies were not cross-connected

Correct Answer: C

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

Which of the following refers to the requirements that dictate when to delete data backups?

- A. Retention policies
- B. Cloud security impact
- C. Off-site storage
- D. Life-cycle management

Correct Answer: A


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

A server administrator has a system requirement to install the virtual OS on bare metal hardware. Which of the following hypervisor virtualization technologies should the administrator use to BEST meet the system requirements? (Choose two.)

- A. Host
- B. Template
- C. Clone
- D. Type 1
- E. Type 2
- F. Guest

Suggested Answer: DF

Community vote distribution



🗲️ 👤 **Obi_Wan_Jacoby** Highly Voted 👍 2 years, 5 months ago

Selected Answer: DF

I believe D & F are accurate. Type 1 hypervisors allow access to the direct hardware (aka bare-metal). Type 2 hypervisors are installed overtop an OS and no longer allow access to direct hardware. The OS the hypervisor is installed on would not be a virtual OS, and this question states it would be a virtual OS being installed on bare-metal. Therefore, it would need to be a "type 1" hypervisor and then "guest" would be the other term as the virtual OS would in-fact be a guest OS.

upvoted 8 times

🗲️ 👤 **Andrewyounan** Highly Voted 👍 2 years, 9 months ago

Selected Answer: DE

Should be D & E since the question is about the hypervisor not the guest OS

upvoted 6 times

🗲️ 👤 **gingasaurusrex** Most Recent 🕒 2 years, 2 months ago

Selected Answer: DF

ChatGPT says this

The two virtualization technologies that would best meet the system requirements for installing a virtual OS on bare metal hardware are:

- D. Type 1
- F. Guest

Type 1 hypervisors, also known as bare-metal hypervisors, are installed directly onto the host server's hardware, and they allow multiple guest operating systems to run on top of them. This provides the closest possible interaction between the virtual machines and the underlying hardware.

Guest refers to the virtual machine or operating system that is running on the hypervisor. By using a type 1 hypervisor, the administrator can create multiple guest VMs running different operating systems on the same physical server, meeting the system requirement to install the virtual OS on bare metal hardware.

upvoted 3 times

🗲️ 👤 **SecNoob27639** 9 months, 2 weeks ago

ChatGPT is a questionably trustworthy source for things like this. I asked it the same question, and it came back with D and A. So while D is pretty much certain, I would recommend using additional study material to determine the correct answer.

upvoted 1 times

A technician is troubleshooting a server issue. Which of the following should the technician do to ensure the solution can be duplicated in the future?

- A. Verify system functionality and implement preventive measures
- B. Document the findings, actions, and outcomes throughout the process
- C. Determine if there is a common element or symptom
- D. Notify the impacted users before implementing any changes

Correct Answer: *B*


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

Which of the following ensures a secondary network path is available if the primary connection fails?

- A. Link aggregation
- B. Most recently used
- C. Heartbeat
- D. Fault tolerance

Suggested Answer: D

Community vote distribution



🗳️ **DannyCarry** 10 months, 3 weeks ago

Link Aggregation combines multiple physical ports into one logical port that does provide redundancy, however the question states "secondary network path". I'm going with D

upvoted 2 times

🗳️ **Fart2023** 11 months ago

Selected Answer: D

Link aggregation also ensures a secondary network path is available if the primary connection fails. Both Link aggregation and Fault tolerance are valid answers, but in the context of specifically ensuring a secondary network path, Link aggregation is the more precise term.

upvoted 1 times

🗳️ **srtysrhtyjumnuyedt** 1 year, 1 month ago

Selected Answer: A

The purpose of link aggregation is to increase bandwidth AND provide fault tolerance.

upvoted 2 times

🗳️ **error77** 1 year, 4 months ago

Selected Answer: A

link agg. includes failover automatically

upvoted 1 times

🗳️ **EngAbood** 1 year, 5 months ago

Selected Answer: D

for sure D:

upvoted 1 times

🗳️ **comptiaboy** 2 years, 2 months ago

Selected Answer: D

Agree with FineB

upvoted 2 times

🗳️ **Fineb** 2 years, 2 months ago

The D is correct. Option A will only Aggregate the link for better performance like bandwidth.

upvoted 3 times

A developer is creating a web application that will contain five web nodes. The developer's main goal is to ensure the application is always available to the end users. Which of the following should the developer use when designing the web application?

- A. Round robin
- B. Link aggregation
- C. Network address translation
- D. Bridged networking

Suggested Answer: A

Community vote distribution

A (100%)

 **gingasaurusrex** 8 months, 3 weeks ago

Selected Answer: A

A. Round robin.

To ensure high availability and distribute the load across multiple web nodes, the developer should use a load balancing technique such as round-robin. With round-robin, incoming requests are distributed evenly across the available web nodes in a circular order, which helps to prevent any one node from becoming overwhelmed with traffic.

Link aggregation, network address translation, and bridged networking are all networking technologies that can be used in various contexts, but they are not specifically designed to ensure high availability of web applications.


upvoted 3 times

Which of the following is the MOST secure method to access servers located in remote branch offices?

- A. Use an MFA out-of-band solution
- B. Use a Telnet connection
- C. Use a password complexity policy
- D. Use a role-based access policy

Suggested Answer: A

Community vote distribution



gingasaurusrex Highly Voted 8 months, 3 weeks ago

Selected Answer: A

A. Use an MFA out-of-band solution.

The most secure method to access servers located in remote branch offices is to use an MFA (multi-factor authentication) out-of-band solution. MFA adds an extra layer of security beyond the traditional username and password by requiring an additional form of authentication, such as a fingerprint or a one-time code sent to a mobile device. An out-of-band solution means that the second factor is sent through a separate channel or device, further reducing the risk of compromise.

Telnet connections are not secure because they transmit data in clear text, making it vulnerable to interception and unauthorized access. Password complexity policies and role-based access policies are important security measures, but they alone are not enough to provide sufficient security for remote access to servers.

upvoted 5 times

King2 Most Recent 1 year, 3 months ago

I think answer is A

upvoted 1 times

King2 1 year, 3 months ago

I think answer is A

MFA for the entire network and also to access a remote server.

Some examples of out-of-band authentication (when paired with a typical online login) include:

Push notifications sent to a mobile device

QR codes with encrypted transaction data

Biometric readers for fingerprint scans or facial recognition

Phone calls for voice authentication

upvoted 1 times

Timock 1 year, 3 months ago

Selected Answer: D

MFA and out of band could be an answer but MFA isn't specific to just secure remote access .. it is for the entire network. Out of band would be authentication that is done outside of the original communication.

Telnet is not secure

A password complexity policy has nothing to do with most secure access method

RBAC provides a consistent authentication and authorization mechanism for users access. Role-based access control (RBAC) restricts network access based on a person's role within an organization and has become one of the main methods for advanced access control. The roles in RBAC refer to the levels of access that employees have to the network.

upvoted 2 times

Fineb 8 months, 1 week ago

A is the correct answer. RBAC without OOB is useless as the server is needed to be accessed remotely

upvoted 1 times

A server administrator is installing a new server that uses 40Gb network connectivity. The administrator needs to find the proper cables to connect the server to the switch. Which of the following connectors should the administrator use?

- A. SFP+
- B. GBIC
- C. SFP
- D. QSFP+

Suggested Answer: D

Community vote distribution

D (100%)

🗨️ 👤 **gingasaurusrex** 8 months, 3 weeks ago

Selected Answer: D

D. QSFP+ is the proper connector to use for 40Gb network connectivity. SFP+ and SFP are connectors commonly used for 10Gb network connectivity, while GBIC is an older connector type that is less commonly used today.

upvoted 1 times

🗨️ 👤 **azre_certified1111** 1 year ago

Selected Answer: D

D is correct.

SFP = 1Gbps

SFP+ = 10Gbps

QSFP+ = 40 Gbps

QSFP28 = 100Gbps

GBIC (Gigabit Interface Converter)

upvoted 3 times

🗨️ 👤 **RSMCT2011** 1 year, 2 months ago

<https://www.3coptics.com/News/28.html>


upvoted 1 times

A technician is laying out a filesystem on a new Linux server. Which of the following tools would work BEST to allow the technician to increase a partition's size in the future without reformatting it?

- A. LVM
- B. DiskPart
- C. fdisk
- D. Format

Suggested Answer: A

Community vote distribution



🗳️ **surfuganda** 9 months, 3 weeks ago

Selected Answer: A

A. LVM

The Logical Volume Manager (LVM) is a device mapper framework that provides logical volume management for the Linux kernel. LVM allows for flexible disk management, enabling the resizing of partitions (logical volumes) on the fly without the need for unmounting or reformatting them. This capability makes it an ideal choice for systems where disk space requirements may change over time, providing administrators with the ability to increase or decrease the size of partitions as needed without disrupting services.

upvoted 2 times

🗳️ **AzadOB** 10 months, 2 weeks ago

Selected Answer: A

LVM why creating confusion ?

upvoted 1 times

🗳️ **H_A_79** 1 year, 3 months ago

LVM is the best tool to allow a technician to increase a partition's size in the future without reformatting it. LVM provides flexibility in managing storage by abstracting physical storage devices into logical volumes that can be resized, extended, or shrunk without the need to reformat or repartition the file system. This makes it an ideal choice for dynamically managing storage on a Linux server.

upvoted 1 times

🗳️ **zoro99** 1 year, 4 months ago

A. LVM

LVM allows a Logical Volume to span multiple physical disks/RAID sets. If you run out of space on your current disk, just add a new disk to the system, use it to extend the volume group that needs more space, and add the necessary space to your logical volume. If you need to move your data to new disks and remove the old ones, you can use pvmove or LVM-level mirroring to migrate the data to the new disks. All this can be done while filesystems are mounted and applications are running.

With FDISK, you can extend an existing partition only if there is free space on the disk immediately after the partition. So extending any partitions other than the last one on the disk will require moving partitions around, which cannot be done while the partition is mounted.

upvoted 2 times

🗳️ **Isaiah44** 1 year, 5 months ago

Selected Answer: B

Diskpart to extend or resize without formatting

upvoted 1 times

🗳️ **comptiaboy** 1 year, 8 months ago

Selected Answer: C

Isn't this C?

upvoted 1 times

Which of the following should be placed at the top of a Bash script to ensure it can be executed?

- A. bash
- B. !execute
- C. #!
- D. @echo off

Suggested Answer: *C*

  **RSMCT2011** 8 months, 1 week ago
<https://earthly.dev/blog/understanding-bash/>
upvoted 2 times


A systems administrator has several different types of hard drives. The administrator is setting up a NAS that will allow end users to see all the drives within the NAS. Which of the following storage types should the administrator use?

- A. RAID array
- B. Serial Attached SCSI
- C. Solid-state drive
- D. Just a bunch of disks

Suggested Answer: D

Community vote distribution

D (100%)



  **surfuganda** 9 months, 3 weeks ago

Selected Answer: D

D. Just a bunch of disks (JBOD)

Just a Bunch Of Disks (JBOD) is a storage architecture that groups multiple hard drives together, allowing each drive to be seen and accessed independently by the operating system or users. This setup does not provide redundancy or performance enhancements like RAID configurations but does enable the use of drives of different sizes and types within the same enclosure or NAS system.

upvoted 2 times

  **Timock** 2 years, 3 months ago

Selected Answer: D

Because of the varying sizes/types of HDs in the question and that it requests end users to see ALL the drives within JBOD would be a better selection than RAID. With a RAID it would appear as one logical unit and the varying sizes would have a limit of the smaller disks used.

<https://www.techtarget.com/searchstorage/definition/JBOD>


upvoted 2 times

Which of the following is the BEST action to perform before applying patches to one of the hosts in a high availability cluster?

- A. Disable the heartbeat network.
- B. Fallback cluster services.
- C. Set the cluster to active-active.
- D. Failover all VMs.

Suggested Answer: D

Community vote distribution



shanebrown 9 months, 2 weeks ago

Selected Answer: D

D. This ensure a noted is actively running while the patch is being applied to another.
upvoted 1 times

surfuganda 1 year, 3 months ago

Selected Answer: D

D. Failover all VMs

This ensures that the VMs remain operational and accessible while the host system is being updated, minimizing downtime and maintaining the availability of services.

Why?

Maintains Service Availability: By failing over VMs to another host, critical services and applications continue to run without interruption.

Minimizes Risk: Moving VMs away from the host being patched reduces the risk of service disruption if unexpected issues arise during the patching process.

Failing over all VMs to another host within the cluster is a direct and effective way to ensure that services remain available and uninterrupted during the maintenance process, adhering to best practices for managing high availability environments.

upvoted 2 times

MrS 1 year, 12 months ago

Selected Answer: D

D. Failover all VMs. This option means to move all the virtual machines running on the host that needs to be patched to another host in the cluster. This way, the host can be safely patched without affecting the availability or performance of the virtual machines. Failover can be done manually or automatically by using a cluster management tool or a hypervisor feature.
upvoted 1 times

Dingos 2 years, 1 month ago

Falling back cluster services is not the BEST action to perform before applying patches to one of the hosts in a high availability cluster because it does not ensure that there is no downtime during the patching process. Failing over all VMs ensures that there is no downtime during the patching process.

upvoted 1 times

Obi_Wan_Jacoby 2 years, 5 months ago

Selected Answer: C

I believe C is the answer. The purpose of Active-Active is to have zero down-time. If a node fails, it automatically sends all that traffic to the other running node (as they were active-active, aka "Live Failover"). This would allow the patching/rebooting to occur with no downtime. Also, it is less work than failing over everything manually etc, and it does state "BEST" action to perform.

upvoted 2 times

Pongsathorn 2 years, 7 months ago

Selected Answer: D

Cluster Patching Process

The process to patch cluster nodes is to bring down one node at a time for patching (and the subsequent reboot) and then return it to service. This process is accomplished for each cluster member, never bringing down all of the nodes simultaneously. It is not usually realistic to manage this

process manually. Services such as Microsoft SCCM and Microsoft Cluster Aware Updating can automate the process.
The Official CompTIA Server+ Study Guide (Exam SK0-005)

I think the most important is the process to patch cluster nodes is to bring down one node at a time for patching (and the subsequent reboot) and then return it to service.

If you have one (or more) VMs on active cluster server, you have to failover the VM(s) to patch the active cluster server.

upvoted 1 times

  **Pongsathorn** 2 years, 7 months ago

Failover and Failback

When you research server clustering, you frequently hear the terms failover and failback. Failover refers to the failure of the active service provider node and passive node's taking over of its responsibilities. At that point, the server that was the passive node becomes the active node. It services clients and handles the cluster's duties. Once the original active node returns to service, you may set it to act as the new passive node, or you may shift the services back to it, causing it to take back its active node role and relegating the formerly passive server to its original role. The term failback refers to the service returning to the original active node.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 242

upvoted 1 times

  **Pongsathorn** 2 years, 7 months ago

I think the "B. Fallback cluster services." is an incorrect spell. Even if they spell correctly it's not the best answer.

If choice B is "Failover cluster services." it will be the best answer to this question.

upvoted 1 times

  **dcdc1000** 2 years, 9 months ago

I also think the answer is D. If you have several vm's on your cluster of servers, you could have; for example, vm1 and vm2 connected to cluster server 1, vm3 connected to cluster server 2 and so on. It makes sense to Failover all VM's located on cluster server 1 to another cluster server prior to applying patches to that server. Just my thoughts.

upvoted 3 times

  **King2** 2 years, 9 months ago

Selected Answer: D

I think the answer is D

upvoted 1 times

  **King2** 2 years, 9 months ago



According to compTIA, Active-Active provides high Availability and load Balancing

Active-Passive provides only high Availability, so the HA cluster mentioned in the question can be Active-Active or Active-Passive, so C is not correct.

Fallback = Role back = alternative which can be used if something goes wrong with the main plan and this is not making sense in this question, so B is not correct.

I think the answer is D.

upvoted 2 times

  **Timock** 2 years, 9 months ago

Secondary servers are already active and waiting to receive connections, so there is no downtime during failover scenario. The environment enjoys improved the processing capacity because two server nodes are actively running instead of one waiting Idly as it does in Active/Passive Clustering.

Therefore you can remove a single instance from the cluster to patch the server and this would be transparent to the users. Failovers may not be.

<https://www.sanitysolutions.com/active-clustering-101/#>

upvoted 1 times

  **jagoichi** 2 years, 9 months ago

B

reference comptia official server + SK0-005



The purpose behind server clusters is to provide high availability of services and data, scalability, and load balancing.

Consider the following example: You have a web site hosting critical customer data.

A BIA determined that if the website becomes unavailable, there will be severe consequences for the organization. As the sysadmin, you decide to create a two-node cluster of web servers. The two servers each host duplicate copies of the files that



make up the web site. If one server fails, the other server is configured to automatically and immediately take over. Since both nodes share the same IP address, there is no interruption or reconfiguration necessary. Also, if you need to take the servers down for maintenance, you can do so one at a time, allowing the site to remain available during patching or upgrade procedures

upvoted 1 times

  **Timock** 2 years, 9 months ago

Two issues with this answer. First it states Fallback and not Failback. Second a failback has to do with the primary server becoming unavailable and then failing back when it is back online. The term here would be failover for what I believe that you are thinking.

upvoted 2 times

  **szl0144** 2 years, 9 months ago

Selected Answer: C

I prefer option C

upvoted 2 times

A server administrator is deploying a new server that has two hard drives on which to install the OS. Which of the following RAID configurations should be used to provide redundancy for the OS?

- A. RAID 0
- B. RAID 1
- C. RAID 5
- D. RAID 6

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **Pongsathorn** 7 months, 3 weeks ago

Selected Answer: B

RAID 1 Disk Mirroring

This RAID design provides a very high degree of redundancy. All content is written to both disks, and if one disk fails, the other disk has everything. This solution is excellent for mission-critical situations. The downside to mirroring is that drive space is not very efficiently used (if you buy two disks that are 1 TB each, you've paid for 2 TB of storage but will only have the effective capacity of 1 TB). RAID 1 requires a minimum of two disks.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 114.

upvoted 1 times

🗨️ 👤 **Timock** 9 months, 2 weeks ago

Selected Answer: B

RAID 1 would be mirroring. RAID 0 striping and therefore no redundancy.

RAIDs 5 & 6 require a minimum 3/4 disks so cannot be the answer here.

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

upvoted 1 times

A systems administrator has been alerted to a zero-day vulnerability that is impacting a service enabled on a server OS. Which of the following would work BEST to limit an attacker from exploiting this vulnerability?

- A. Installing the latest patches
- B. Closing open ports
- C. Enabling antivirus protection
- D. Enabling a NIDS

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **a792193** 1 year ago

Selected Answer: B

Correct Answer: B. Closing open ports

Incorrect answer: NIDS - because the "D" stands for detection, not PREVENTION. NIDS would not stop an attack it would only generate an alert.
upvoted 1 times

🗳️ 👤 **hasquaati** 1 year, 3 months ago

Selected Answer: B

- Nothing in the question is asking to detect the zero day attack, so NIDS is not relevant. We already know about the attack.
 - Nothing in the question is asking us to keep the service running with our action.
 - There is a low probability that a patch will solve an issue of a zero day attack.
 - Closing the port is the only option that can prevent malware traffic effecting the particular service.
- upvoted 1 times

🗳️ 👤 **MrS** 1 year, 5 months ago

Selected Answer: B

B. Closing open ports. This option means blocking or disabling the network ports used by the service affected by the zero-day vulnerability. Closing open ports can work best to limit an attacker from exploiting the zero-day vulnerability, as it can prevent or reduce the exposure of the service to the network and reduce the attack surface. Closing open ports can also help to isolate the host from potential attacks and minimize the impact on other hosts or systems in the cluster.
upvoted 1 times

🗳️ 👤 **RSMCT2011** 2 years, 2 months ago

assumption: a zero-day vulnerability means the patch is still not available, so answer is D: Enabling a NIDS
upvoted 2 times

🗳️ 👤 **RSMCT2011** 2 years, 2 months ago

Since we know the services with zero-day vulnerabilities, I think closing ports is better than installing NIDS. so B is a better answer
upvoted 2 times

🗳️ 👤 **Drewid91** 2 years, 2 months ago

Depending on business needs, closing a port may not be a viable option. NIDS feels like the most likely answer to be correct.
upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 1 year, 11 months ago

But is not NIDS a "detection" system? The questions asks what is best to limit an attacker from exploiting the vulnerability, whereas a NIDS will simply tell you once it happened. In this case, answer B would be it (Assuming they are closing the specific ports utilized in the zero-day). A&C obviously are no good against zero-day.
upvoted 2 times

A technician is checking a server rack. Upon entering the room, the technician notices the fans on a particular server in the rack are running at high speeds. This is the only server in the rack that is experiencing this behavior. The ambient temperature in the room appears to be normal. Which of the following is the MOST likely reason why the fans in that server are operating at full speed?

- A. The server is in the process of shutting down, so fan speed operations have been defaulted to high.
- B. An incorrect fan size was inserted into the server, and the server has had to increase the fan speed to compensate.
- C. A fan failure has occurred, and the other fans have increased speed to compensate.
- D. The server is utilizing more memory than the other servers, so it has increased the fans to compensate.

Suggested Answer: C

Community vote distribution

C (100%)

 **gingasaurusrex** 8 months, 3 weeks ago

Selected Answer: C

C. A fan failure is the MOST likely reason why the fans in that server are operating at full speed while the other servers in the rack are operating normally. When a fan fails, the other fans in the server have to work harder to compensate for the loss of airflow. This increased workload causes the remaining fans to run at a higher speed than normal to maintain adequate cooling for the server components.

Option A is incorrect because the fan speed is not dependent on the server shutting down. Option B is unlikely because the server would not be designed to increase the fan speed to compensate for an incorrect fan size. Option D is also unlikely because fan speed is not dependent on the memory usage of the server.


upvoted 3 times

A server administrator is installing an OS on a new server. Company policy states no one is to log in directly to the server. Which of the following installation methods is BEST suited to meet the company policy?

- A. GUI
- B. Core
- C. Virtualized
- D. Clone

Suggested Answer: B

Community vote distribution



🗳️ 👤 **Evanj51** Highly Voted 2 years, 1 month ago

Shouldn't this be virtualized
upvoted 6 times

🗳️ 👤 **Obi_Wan_Jacoby** 1 year, 11 months ago

I agree, if you install Core onto a new server, you can still log into it via console
upvoted 3 times

🗳️ 👤 **cj207800** Most Recent 5 months, 1 week ago

Selected Answer: B

This is what I found. A core installation, also known as a headless or minimal installation, provides a command-line interface without a graphical user interface (GUI). This approach aligns with the company policy for several reasons:

Reduced attack surface: A core installation eliminates unnecessary components and services, minimizing potential vulnerabilities.

Remote management: Core installations are typically managed remotely through secure protocols like SSH, eliminating the need for direct local logins.

Efficiency: Core installations use fewer system resources, allowing the server to dedicate more power to its primary functions.

Compliance: By not having a GUI, it discourages direct logins and encourages proper remote management practices, adhering to the company policy.
upvoted 1 times

🗳️ 👤 **maigoya** 9 months, 2 weeks ago

Selected Answer: C

A. GUI (Graphical User Interface): Installing the OS through a graphical interface typically requires physically connecting a monitor, keyboard, and mouse to the server. This would involve direct login access, potentially violating the company policy.

B. Core: While a core installation might reduce the server's attack surface compared to a full GUI installation, it wouldn't eliminate the need for direct login for configuration.

D. Clone: Cloning an existing server creates a copy of the existing system configuration. While potentially faster for initial setup, it wouldn't address the core policy of avoiding direct login to the new server, especially if the cloned server contains unnecessary accounts or security vulnerabilities.
upvoted 2 times

🗳️ 👤 **gingasaurusrex** 1 year, 8 months ago

Selected Answer: B

Chat GPT said this



B. Core installation method is best suited to meet the company policy.

A core installation method installs only the essential components required to run the operating system. It does not include a graphical user interface (GUI) or any other extra features that are not required for the server's function. This minimal installation reduces the attack surface of the server, making it more secure.

Since company policy states that no one is to log in directly to the server, a core installation is the best choice as it removes any unnecessary components and services that could be exploited by unauthorized users.

Virtualized installation and clone installation are not relevant to the company policy of prohibiting direct logins. A GUI installation includes unnecessary features and services that could pose security risks. Therefore, a core installation is the best option for meeting the company policy.

upvoted 1 times

  **kloug** 1 year, 9 months ago

Bbbbbbbbbb

upvoted 1 times

Which of the following licenses would MOST likely include vendor assistance?


- A. Open-source
- B. Version compatibility
- C. Subscription
- D. Maintenance and support

Suggested Answer: C

Community vote distribution

C (89%)

11%

 **shanebrown** 9 months, 2 weeks ago

Selected Answer: C

Subscriptions normally applies to software while maintenance is normally applied to hardware. Additionally, subscription is a type of license while maintenance is an agreement/contract. I hope this helps.

upvoted 2 times

 **Grumpy_Old_Coot** 1 year, 8 months ago

Selected Answer: D

Wording. Maintenance and Support contracts don't normally support server software, but server hardware.

upvoted 1 times

 **Pongsathorn** 2 years, 7 months ago

Selected Answer: C

Maintenance and Support Plans

It is important not to confuse the licensing and maintenance concepts. Just because you have the legal right to use a piece of software or install an OS does not necessarily entitle you to vendor support for the product. Some licensing models may include support. Subscription license models are an example. Support plans are covered later in the book.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 22.

upvoted 2 times

 **nixonbii** 2 years, 8 months ago

Selected Answer: C

I agree. I have never heard of a "maintenance and support" license.


upvoted 1 times

 **Andrewyounan** 2 years, 9 months ago

Selected Answer: C

Subscription - is type of license that can be supported and maintained by vendor

upvoted 2 times

 **Timock** 2 years, 9 months ago

Selected Answer: C

Subscription - Maintenance and support are normally a part of a license and not listed as a license type itself - Neither is version compatibility. Open-source is free and doesn't normally offer any kind of support except communities.

upvoted 2 times

A company wants to deploy software to all users, but very few of them will be using the software at any one point in time. Which of the following licensing models would be BEST for the company?

- A. Per site
- B. Per concurrent user
- C. Per core
- D. Per instance

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **Pongsathorn** 7 months, 3 weeks ago

Selected Answer: B

Per-concurrent-user—one license for each software instance in use by a user. This is typically less expensive than per-seat licensing. If your organization has ten of these licenses, and there are twenty copies of the software installed, then only ten users may use the software simultaneously.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 23.

upvoted 3 times

A technician is setting up a physical server that will be used to store sensitive data. The technician plans to set up the internal hard drives in a RAID 5 array for redundancy. Which of the following would be the BEST type of interface for the technician to use?

- A. SAS
- B. USB
- C. SD
- D. eSATA

Suggested Answer: A

Community vote distribution

A (100%)

 **gingasaurusrex** 8 months, 3 weeks ago

Selected Answer: A

A. SAS (Serial Attached SCSI) would be the BEST type of interface for the technician to use for setting up a RAID 5 array for storing sensitive data on a physical server. SAS provides higher performance, reliability, and data transfer rates compared to other interface types such as USB, SD, or eSATA. Additionally, SAS has advanced features like dual-porting, which provides redundancy and multipathing capabilities for improved reliability and availability of data. Therefore, SAS is a popular choice for enterprise-level storage applications that require high-performance and reliable storage solutions.

upvoted 2 times

Which of the following should the technician attempt FIRST to resolve the issue?

- A. mkdir /home
B. umount nfssserver:/home
C. rmdir nfssserver:/home/dev/sda2
D. mount /home

Community vote distribution



Selected Answer: D

"B. umount nfsserver:/home" - This command will result in an error because the NFS mount has not been done.

Therefore, executing “D. mount /home” is the first and most appropriate step to solve this problem. This will hopefully result in a correct NFS mount based on the settings in /etc/fstab and make the user’s files available again.

upvoted 1 times

Selected Answer: D

The `mount /home` command attempts to mount `/home` based on the existing entry in `/etc/fstab`. This would attempt to mount `nfserver:/home` to `/home` as specified. Since the issue seems to be that the NFS share did not mount properly, this is the correct first step to attempt to resolve the issue.

upvoted 1 times

Selected Answer: B

B. Mount the NFS share using the mount command. This option means using the command-line tool to connect the remote NFS share to the local mount point. Mounting the NFS share can resolve the issue, as it can restore access to the files in the home directory that are stored on the NFS server. Mounting the NFS share can also verify that the NFS client package is installed and configured correctly on the Linux server.

upvoted 2 times

Except B isn't mounting, it's unmounting: "umount"

upvoted 1 times

Bbbbbbbbbbbb

upvoted 1 times

D: is the first thing to try. A correct fstab entry and /home folder exists..

upvoted 1 times

Selected Answer: A

The fstab (/etc/fstab) (or file systems table) file is a system configuration file on Debian systems. The fstab file typically lists all available disks and disk partitions, and indicates how they are to be initialized or otherwise integrated into the overall system's file system.

The etc/fstab file has /home as the DIR.. but when a df /home is asked for there is no /home directory underneath... only /dev/sda2. Therefore, we would need to mkdir a /home under the /home directory.

```
# <file system> <dir> <type> <options> <dump> <pass>
10.10.0.10:/backups /var/backups nfs defaults 0 0
```

<https://linuxize.com/post/how-to-mount-an-nfs-share-in-linux/>
upvoted 1 times

  **Pongsathorn** 2 years, 7 months ago

Run the mount command in one of the following forms to mount the NFS share:

```
mount /var/backups
mount 10.10.0.10:/backups
```

Copy

The mount command, will read the content of the /etc/fstab and mount the share.

Next time you reboot the system the NFS share will be mounted automatically.

<https://linuxize.com/post/how-to-mount-an-nfs-share-in-linux/>
upvoted 1 times

  **lordguck** 2 years, 7 months ago

A is wrong as /home exists! if you use df -h /blah and the folder does not exists, you get a "directory not found" reply


upvoted 1 times

A company is reviewing options for its current disaster recovery plan and potential changes to it. The security team will not allow customer data to egress to non-company equipment, and the company has requested recovery in the shortest possible time. Which of the following will BEST meet these goals?

- A. A warm site
- B. A hot site
- C. Cloud recovery
- D. A cold site

Suggested Answer: B

Community vote distribution



🗳️ 👤 **2e6ef4a** 3 weeks, 6 days ago

Selected Answer: B

B. Duh

upvoted 1 times

🗳️ 👤 **Kraken84** 8 months, 2 weeks ago

Given the constraints mentioned:

Customer data should not egress to non-company equipment.

Recovery in the shortest possible time.

The best option would be:

B. A hot site

Given the requirements, a hot site is the best choice because it meets the need for rapid recovery and can be owned by the company, ensuring customer data does not egress to non-company equipment.

upvoted 1 times

🗳️ 👤 **kloug** 1 year, 3 months ago

Bbbbbbbbbb

upvoted 1 times

🗳️ 👤 **Obi_Wan_Jacoby** 1 year, 5 months ago

Selected Answer: B

Answer is B. Companies may choose to build/maintain their own hot site. And the questions asks for recovery in the shortest possible time.

upvoted 3 times

🗳️ 👤 **lordguck** 1 year, 7 months ago

B: is correct. It's a replication of the primary DC in all aspects with data replication. There is no rule, that the equipment must belong to another company (see comment of Timock2)

upvoted 2 times

🗳️ 👤 **nixonbii** 1 year, 8 months ago

What if the company is Amazon or Microsoft? Given the AWS and MS Azure services they provide, it seems as though they actually own their hot sites.

upvoted 1 times

🗳️ 👤 **Timock** 1 year, 9 months ago

Selected Answer: A

A warm site does not have the data and would fit with the no data on a NON-company equipment requirement.

A hot side and cloud recovery would have non-company equipment and the cold site is not the shortest possible time.

upvoted 1 times

Which of the following encryption methodologies would MOST likely be used to ensure encrypted data cannot be retrieved if a device is stolen?

- A. End-to-end encryption
- B. Encryption in transit
- C. Encryption at rest
- D. Public key encryption

Suggested Answer: C

Community vote distribution

C (75%)

B (25%)

🗳️ **surfuganda** 9 months, 3 weeks ago

Selected Answer: C

C. Encryption at rest

Encryption at rest is specifically designed to protect data stored on devices such as hard drives, USB drives, or mobile devices. When data is encrypted at rest, it remains encrypted even if the device is stolen or compromised. This ensures that the data cannot be accessed without the appropriate decryption keys, thereby safeguarding it from unauthorized access.

upvoted 2 times

🗳️ **Frank9020** 11 months, 2 weeks ago

Selected Answer: C

C. Encryption at rest

Encryption at rest is a method used to encrypt data stored on a device or media, such as hard drives, ensuring that even if the device is stolen, the data remains secure. This encryption protects the data when it is not actively being used, making it a suitable choice for scenarios where the primary concern is securing data on storage devices.

upvoted 1 times

🗳️ **RBL23168** 1 year, 1 month ago

C. Encryption at rest

Encryption at rest is specifically designed to protect data stored on devices such as laptops, mobile devices, and servers. If a device is stolen, encryption at rest ensures that the data remains encrypted and inaccessible to unauthorized individuals. This type of encryption is applied to the storage media itself, making it difficult for someone to retrieve sensitive information even if they have physical possession of the device.

upvoted 1 times

🗳️ **RBL23168** 1 year, 1 month ago

C. Encryption at rest

Encrypt it while you still have control. Once it's stolen, it's unusable.

upvoted 1 times

🗳️ **Sjabs** 1 year, 2 months ago

Selected Answer: B

answer is B

upvoted 1 times

A company stores extremely sensitive data on an air-gapped system. Which of the following can be implemented to increase security against a potential insider threat?

- A. Two-person integrity
- B. SSO
- C. SIEM
- D. Faraday cage
- E. MFA

Suggested Answer: A

Community vote distribution

A (100%)

🗨️ 👤 **Pongsathorn** 7 months, 1 week ago

Selected Answer: A

An air gap, air wall, air gapping[1] or disconnected network is a network security measure employed on one or more computers to ensure that a secure computer network is physically isolated from unsecured networks, such as the public Internet or an unsecured local area network.[2] It means a computer or network has no network interface controllers connected to other networks,[3][4] with a physical or conceptual air gap, analogous to the air gap used in plumbing to maintain water quality.

upvoted 1 times

🗨️ 👤 **RSMCT2011** 8 months, 1 week ago

https://csrc.nist.gov/glossary/term/two_person_integrity

upvoted 1 times

Which of the following can be used to map a network drive to a user profile?

- A. System service
- B. Network service
- C. Login script
- D. Kickstart script

Suggested Answer: C

Community vote distribution

C (100%)


  **Pongsathorn** 7 months, 3 weeks ago

Selected Answer: C

A login script is a script that is executed when a user logs into a computer. A login script can adjust settings in the operating system, map network drives for different groups of users, or even display a welcome message that is specific to each user. Multiple login scripts can even be utilized at the same time, with specific ones activated based on which user logs in and the operating system in use on the computer.

<https://www.computerhope.com/jargon/l/login-script.htm>

upvoted 3 times

  **Pongsathorn** 7 months, 3 weeks ago

Login and Logout Scripts

Some scripts execute based on the user. For example, a login script runs when a user logs in. The script customizes the user's environment, maps network drives, and automatically launches a particular application. Logout scripts execute when the user signs off. These scripts might delete temp files or copy the user's data to a network server.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 160.


upvoted 2 times

A server shut down after an extended power outage. When power was restored, the system failed to start. A few seconds into booting, the Num Lock, Scroll Lock, and Caps Lock LEDs flashed several times, and the system stopped. Which of the following is the MOST likely cause of the issue?

- A. The keyboard is defective and needs to be replaced.
- B. The system failed before the display card initialized.
- C. The power supply is faulty and is shutting down the system.
- D. The NIC has failed, and the system cannot make a network connection.

Suggested Answer: B

Community vote distribution



🗳️ 👤 **Kabbath1986** 2 months, 1 week ago

Selected Answer: C

✓ C. The power supply is faulty and is shutting down the system.

Here's why:

The flashing of the Num Lock, Scroll Lock, and Caps Lock LEDs is a classic POST error code (Power-On Self-Test). This often indicates a hardware-level failure, and when it happens early during boot, it's usually related to critical system components like the power supply, motherboard, or memory.

upvoted 1 times

🗳️ 👤 **shanebrown** 9 months, 2 weeks ago

Selected Answer: C

C. I'll go with C because it states there was a power outage before the other information was presented about the keyboard lights. Which is pointing to PSU. There could have been a power spike that could affect the PSU.

upvoted 1 times

🗳️ 👤 **badgerino** 1 year, 2 months ago

Selected Answer: C

How do we know it's the video card or display card? It doesn't specifically state if we're able to see the POST/BIOS screen.. not enough information it could be either B or C. The lights on the keyboard flashing does not indicate it's a video card issue to me. It could be any component like motherboard, CPU, RAM as surfuganda stated. PSU being faulty, under load and tripping makes more sense to me in this scenario.

upvoted 1 times

🗳️ 👤 **surfuganda** 1 year, 3 months ago

Selected Answer: B

B. The system failed before the display card initialized.

Explanation:

When a system fails to boot and the Num Lock, Scroll Lock, and Caps Lock LEDs flash, it usually indicates a hardware failure. However, the fact that the failure occurs before the display card initializes suggests that the issue lies deeper in the system, likely with core hardware components such as the CPU, motherboard, or RAM.

upvoted 1 times

🗳️ 👤 **K1lroy** 1 year, 11 months ago

Selected Answer: C

Chatgpt agrees with C:

"A problematic Power Supply Unit is more likely to cause the described behavior. An issue with the PSU could lead to unstable power delivery, which may trigger the server's built-in protection mechanisms, causing it to shut down. Faulty power supply units are known to cause erratic behavior, including unusual light patterns on the indicator LEDs."

upvoted 1 times

🗳️ 👤 **MrS** 1 year, 12 months ago

Selected Answer: B

The MOST likely cause of the issue is that the system failed before the display card initialized. The flashing of the Num Lock, Scroll Lock, and Caps Lock LEDs is a BIOS error code indicating that the system has failed before the display card is initialized.


upvoted 3 times

A systems administrator is investigating a server with a RAID array that will not boot into the OS. The administrator notices all the hard drives are reporting to be offline. The administrator checks the RAID controller and verifies the configuration is correct. The administrator then replaces one of the drives with a known-good drive, but it appears to be unavailable as well. Next, the administrator takes a drive out of the server and places it in a spare server, and the drive is available and functional. Which of the following is MOST likely causing the issue?

- A. The kernel is corrupt.
- B. Resources are misallocated.
- C. The backplane has failed.
- D. The drives need to be reseated.

Suggested Answer: C

Community vote distribution



Pongsathorn Highly Voted 1 year, 1 month ago

Selected Answer: C

There may be many causes of storage problems on a server or its related components.

Boot errors might cause by:

Misconfigured RAID array

Drive failure

Controller or HBA failure

Loose connections or cable failures

Corrupt boot sector

Corrupt file system table

Backplane failure

The server's backplane provides a communications bus among devices that you may install. The backplane is an additional circuit board attached to the motherboard, and not all servers have one. Damage to the backplane—either physical or electrical—may keep devices on this bus from communicating correctly.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 122.

upvoted 5 times

Katlegobogosi Most Recent 10 months, 1 week ago

Selected Answer: C

C. The backplane has failed.

upvoted 2 times

Spacecluster 1 year, 2 months ago

Selected Answer: A

A looks correct answer

upvoted 1 times

Which of the following techniques can be configured on a server for network redundancy?

- A. Clustering
- B. Virtualizing
- C. Cloning
- D. Teaming

Suggested Answer: D

Community vote distribution

D (100%)

🗨️ **msspeedy** 2 months, 1 week ago

Selected Answer: D

Key work network redundancy.

Teaming is about aggregating network connections for better performance and resilience.

Clustering is about grouping servers for high availability and capacity.

upvoted 1 times

🗨️ **PaytoPlay** 4 months, 1 week ago

Selected Answer: A

If I'm running two firewalls as a cluster in active-passive wouldn't this be considered redundancy for the network?

upvoted 1 times

🗨️ **Frank9020** 11 months, 2 weeks ago

Selected Answer: D

D. Teaming

Network teaming, also known as NIC teaming or bonding, involves grouping multiple network interfaces together to form a single, logical network interface. This configuration provides redundancy and load balancing, enhancing network reliability. If one network interface fails, the traffic is automatically routed through the remaining interfaces, ensuring continued network connectivity. Teaming can be configured to operate in various modes, including fault tolerance for redundancy and load balancing for improved performance.

upvoted 1 times

🗨️ **David_IT_guy** 11 months, 4 weeks ago

NOT D.

Teaming (or more loosely, bridging) makes multiple network cards appear as one card. So for example if you have 2 NICs at 1Gb speed, you can team them together to get up to 2Gb of throughput. With teaming, it is very important to make sure that the network hardware you're plugging into supports it.


upvoted 1 times

A server has experienced several component failures. To minimize downtime, the server administrator wants to replace the components while the server is running. Which of the following can MOST likely be swapped out while the server is still running? (Choose two.)

- A. The power supply
- B. The CPU
- C. The hard drive
- D. The GPU
- E. The cache
- F. The RAM

Suggested Answer: AC

Community vote distribution



error77 9 months, 2 weeks ago

Selected Answer: AC

24/7 server rarely runs on a single power supply and single harddisk
upvoted 2 times

surfuganda 9 months, 3 weeks ago

Selected Answer: AC

Forgot to vote
upvoted 1 times

surfuganda 9 months, 3 weeks ago

A. The power supply
and
C. The hard drive

You are outside your mind if you think Hot-pluggable RAM is MORE LIKELY than Hot-pluggable hard drive.

upvoted 2 times

H_A_79 1 year, 3 months ago

A and C are the correct Answer
upvoted 2 times

jjwelch00 1 year, 7 months ago

Selected Answer: AF

A. The power supply - In many modern servers, power supplies are designed to be hot-swappable, meaning they can be replaced without powering down the server.

F. The RAM - In some servers, RAM can be added or removed without shutting down the server, especially if the server has redundant memory modules and supports memory hot-plugging.

upvoted 1 times

jjwelch00 1 year, 7 months ago

The components that can MOST likely be swapped out while the server is still running are:

A. The power supply - In many modern servers, power supplies are designed to be hot-swappable, meaning they can be replaced without powering down the server.

F. The RAM - In some servers, RAM can be added or removed without shutting down the server, especially if the server has redundant memory modules and supports memory hot-plugging.

The other components listed (CPU, hard drive, GPU, and cache) are typically not hot-swappable, and require the server to be powered down before they can be replaced.

upvoted 1 times

An administrator is configuring a host-based firewall for a server. The server needs to allow SSH, FTP, and LDAP traffic. Which of the following ports must be configured so this traffic will be allowed? (Choose three.)

- A. 21
- B. 22
- C. 53
- D. 67
- E. 69
- F. 110
- G. 123
- H. 389

Suggested Answer: ABH

Community vote distribution

ABH (100%)

🗳️ 👤 **surfuganda** 9 months, 3 weeks ago

Selected Answer: ABH

- A. 21 (FTP)
 - B. 22 (SSH)
 - H. 389 (LDAP)
- upvoted 1 times

🗳️ 👤 **gingasaurusrex** 1 year, 8 months ago

Selected Answer: ABH

- B. 22 (SSH)
- A. 21 (FTP)
- H. 389 (LDAP)

Port 53 is used for DNS

Ports 67 and 69 are used for DHCP

Port 110 is used for POP3 email retrieval

Port 123 is used for NTP (network time protocol)

upvoted 1 times

🗳️ 👤 **mdeckard** 1 year, 9 months ago

Selected Answer: ABH

- Port 21 - FTP
 - Port 22 - SSH
 - Port 389 - LDAP
- upvoted 1 times

A technician needs to configure a server's RAID array for maximum capacity. Which of the following RAID levels BEST meets this requirement?

- A. 0
- B. 1
- C. 5
- D. 6

Suggested Answer: A

  **ccoli** 11 months, 2 weeks ago

in 20 years of doing IT work I have never encountered a scenario where RAID 0 was an acceptable solution. Questions like this are why no one respects comptia certs.

upvoted 2 times

A user is unable to modify files on a Windows fileshare. A technician examines the system and determines the user is a member of Group A. Group A was granted full control over the files. Which of the following is MOST likely causing the access issues?

- A. SIEM policies
- B. IDS rules
- C. Share permissions
- D. Firewall ports

Suggested Answer: C

Community vote distribution

C (100%)

 **gingasaurusrex** 8 months, 3 weeks ago

Selected Answer: C

C. Share permissions

The most likely cause of the access issue is the share permissions. Even though Group A was granted full control over the files, it's possible that the share permissions are restricting access to the files. Share permissions are used to control access to a shared folder or file across the network, and they can be set independently of the NTFS permissions that control access to the files on the local system.

SIEM policies, IDS rules, and firewall ports are all security measures that could potentially impact access to the files, but they are less likely to be the cause of the access issue in this scenario.

upvoted 1 times

 **Pongsathorn** 1 year, 1 month ago

Selected Answer: C

Share Permissions

When a resource is shared over a network, it can be assigned Read, Change, Full Control, or Deny permissions on the share. Windows supports user-level security; users have permissions depending upon who they are. Share permissions only protect the resource from access over the network.

Local access can only be secured by NTFS permissions.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 264.

Resource Access

Many help desk tickets and user complaints center around access to resources. Frequently, these issues are tied to permissions. In Windows, file access is managed by NTFS permissions and by Share permissions when accessing the file from across the network. In Linux, file access is controlled by standard permissions and by NFS permissions when accessing the data from across the network. Ensuring these permissions are set correctly eases many user frustrations and provides users access to the files that they need. Viruses and other malware also impact the ability of users to access resources.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 224.

upvoted 3 times

Users are experiencing issues when trying to access resources on multiple servers. The servers are virtual and run on an ESX server. A systems administrator is investigating but is unable to connect to any of the virtual servers. When the administrator connects to the host, a purple screen with white letters appears. Which of the following troubleshooting steps should the administrator perform FIRST?

- A. Check the power supplies.
- B. Review the log files.
- C. Reinstall the ESX server.
- D. Reseat the processors.

Suggested Answer: B

Community vote distribution

B (100%)

  **gingasaurusrex** 8 months, 3 weeks ago

Selected Answer: B

B. Review the log files.

The purple screen with white letters is commonly known as the "purple screen of death" (PSOD), which is an error screen that appears when there is a fatal error on an ESX/ESXi host. The first step in troubleshooting this issue should be to review the log files, which can provide additional information about the error that occurred.

Checking the power supplies or reseating the processors would be unlikely to resolve this type of issue, as it is typically a software or configuration issue rather than a hardware issue. Reinstalling the ESX server would be a drastic step that should only be taken if all other troubleshooting steps have failed.

upvoted 2 times

  **Pongsathorn** 1 year, 1 month ago



Selected Answer: B

Memory Failures

Memory sticks are particularly susceptible to power fluctuations and static electricity. Failed or failing RAM causes OS instability. With Windows Server, this usually causes a Blue Screen to occur. This Windows crash screen provides some clues as to the problem on screen or via a memory dump file that can help identify problems leading up to the crash. If the server successfully restarts, check the log files for more information. On VMware ESXi systems, memory errors may trigger a Purple crash screen. On Linux servers, a kernel panic error may indicate problems with RAM.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 102.

upvoted 2 times

  **Timock** 1 year, 3 months ago

Selected Answer: B

A purple screen of death (PSOD) is a diagnostic screen with white type on a purple background that's displayed when the VMkernel of a VMware ESXi host experiences a critical error, becomes inoperative and terminates any virtual machines (VMs) that are running.

The other answers could be the issue but are a bit drastic and definitely not something you do FIRST.

<https://www.techtarget.com/searchvmware/definition/Purple-Screen-of-Death-PSOD>

upvoted 2 times

Joe, a technician, wants to configure a server's networking information so he will no longer need to maintain a list of names and IP addresses in a file on the server. Which of the following will work BEST to accomplish this task?

- A. DHCP
- B. Hosts file
- C. DNS
- D. VLAN

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Grumpy_Old_Coot** 7 months, 3 weeks ago

Selected Answer: C

C. "List of names and ip#s in file on the computer" is a hosts file. The next step up is DNS. Using a DHCP doesn't do -anything- without having a DNS for it to talk to.

upvoted 1 times

🗳️ 👤 **gingasaurusrex** 1 year, 2 months ago

Selected Answer: C

C. DNS

Domain Name System (DNS) is the best solution to configure a server's networking information so that a list of names and IP addresses will no longer need to be maintained in a file on the server. DNS is a distributed system that translates domain names to IP addresses and vice versa. By using DNS, servers can be identified by their domain names instead of their IP addresses, and the IP addresses can be managed by the DNS server.

DHCP can assign IP addresses to servers dynamically, but it does not eliminate the need for maintaining a list of names and IP addresses. A hosts file can be used to map names to IP addresses, but it needs to be maintained manually on each server. VLAN is a technology used to separate traffic on a network, but it is not related to managing server names and IP addresses.

upvoted 2 times

🗳️ 👤 **agabeen** 1 year, 4 months ago

so sad to see A is the answer ..

C definitely is the Correct one

upvoted 1 times

🗳️ 👤 **Pongsathorn** 1 year, 7 months ago

Selected Answer: C

DNS resolves "names and IP address" things

upvoted 3 times

🗳️ 👤 **nixonbii** 1 year, 8 months ago

Selected Answer: C

So sad. If they had just left out the reference to names, the DHCP answer would have won. Once you map IP addresses to names in a table, it's all over. DNS wins.

upvoted 2 times

🗳️ 👤 **szl0144** 1 year, 9 months ago

Selected Answer: C

DNS seems correct. to replace the host zone file with a list of ip address and host name

upvoted 2 times

A server administrator added a new drive to a server. However, the drive is not showing up as available. Which of the following does the administrator need to do to make the drive available?

- A. Partition the drive.
- B. Create a new disk quota.
- C. Configure the drive as dynamic.
- D. Set the compression.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **MrS** 12 months ago

Selected Answer: A

The administrator needs to partition the drive to make it available. Partitioning is the process of dividing a single hard drive into multiple logical drives. Once the drive is partitioned, it will be available for use.

upvoted 2 times

🗳️ 👤 **gingasaurusrex** 1 year, 2 months ago

Selected Answer: A

A. Partition the drive.

The administrator needs to partition the drive to make it available for use. Partitioning involves dividing the disk into one or more sections that can be used for storing data. Once the partition is created, the administrator can format it with a file system and assign a drive letter or mount point to make it accessible.

Creating a new disk quota, configuring the drive as dynamic, or setting the compression are not necessary steps to make the drive available. Disk quotas are used to limit the amount of disk space that can be used by a user or group. Dynamic disks are used for advanced disk configurations, such as creating spanned or striped volumes. Compression is a feature that can be used to reduce the size of files on a disk.

upvoted 1 times

🗳️ 👤 **nixonbii** 1 year, 8 months ago

Selected Answer: A

Agreed.

upvoted 1 times

🗳️ 👤 **Timock** 1 year, 9 months ago

Selected Answer: A

Disk Management displays whether a disk is online (available), or offline.

In Windows, by default, all newly-discovered disks are brought online with read and write access. In Windows Server, by default, newly-discovered disks are brought online with read and write access unless they are on a shared bus (such as SCSI, iSCSI, Serial Attached SCSI, or Fibre Channel). Disks on a shared bus are offline the first time they are detected.

You would first need to bring online (available) then choose the partition and during this choose simple or spanned... etc. BCD are just options that you can choose later.

<https://learn.microsoft.com/en-us/windows-server/storage/disk-management/manage-disks>

upvoted 2 times

Which of the following are measures that should be taken when a data breach occurs? (Choose two.)


- A. Restore the data from backup.
- B. Disclose the incident.
- C. Disable unnecessary ports.
- D. Run an antivirus scan.
- E. Identify the exploited vulnerability.
- F. Move the data to a different location.

Suggested Answer: BE

Community vote distribution

BE (75%)

CE (25%)

 **surfuganda** 9 months, 3 weeks ago

Selected Answer: BE

- B. Disclose the incident.
 - E. Identify the exploited vulnerability.
- upvoted 1 times

 **error77** 10 months, 2 weeks ago

Selected Answer: BE

- C is wrong - vulnerability may come from a necessary port.
- upvoted 1 times

 **gingasaurusrex** 1 year, 8 months ago

Selected Answer: BE

- B. Disclose the incident.
- E. Identify the exploited vulnerability.

When a data breach occurs, two important measures that should be taken are to disclose the incident and identify the exploited vulnerability. Disclosing the incident is important for transparency and to allow affected individuals to take steps to protect themselves. Identifying the exploited vulnerability can help to prevent future breaches and strengthen security measures.

Restoring the data from backup may be necessary if the data was lost or corrupted during the breach, but it is not a measure that should be taken in all cases. Disabling unnecessary ports, running an antivirus scan, and moving the data to a different location may be necessary steps to prevent further damage or to secure the environment, but they are not measures that should be taken immediately following a data breach.

upvoted 1 times

 **Pongsathorn** 2 years, 1 month ago

Selected Answer: BE

Answer should be B and E.
Regarding CompTIA Server+ SK0-005 Objective. Objective 3.4 Explain data security risks and mitigation strategies.
It's obvious they mentioned "identification" and "disclosure".

upvoted 1 times

 **Pongsathorn** 2 years, 1 month ago

IDENTIFICATION


First, the organization must know when it has been breached. It might surprise you to know that in many cases organizations don't even know that a breach has occurred for weeks or months! Identifying breaches involves deep inspection of log files by experienced technicians to identify that data loss has occurred.

DISCLOSURE

In many highly regulated industries, organizations are required by regulation or law to notify any users whose data has been disclosed. For example, in the healthcare field, the HIPAA Breach Notification Rule, 45 CFR §§ 164.400-414, requires HIPAA-covered entities and their business associates to provide notification following a breach of unsecured protected health information (PHI). As another example, all 50 states, the

District of Columbia, Guam, Puerto Rico, and the Virgin Islands have enacted legislation requiring private or governmental entities to notify individuals of security breaches of information involving personally identifiable information (PII).



upvoted 1 times

  **nixonbii** 2 years, 2 months ago

Selected Answer: BE

The Department of Justice takes a very dim view towards those who lose customers' personal information and then wait to tell the authorities. I don't care if this one gets marked wrong on the exam, if I ever find myself in the middle of such a situation, I'm going to sing like a bird.

upvoted 2 times

  **Timock** 2 years, 3 months ago

Selected Answer: CE

Restoring from backup is unnecessary here as we don't know what exactly the breach affected.

Disclose incident is AFTER all other steps have been taken. The question states WHEN a data breach occurs.

Move the data to another location ... same issue. What data. Breach scope needs to be identified.

Identify the exploited vulnerability -- definitely

Disable ports and running antivirus scan between these two... antivirus should already have been running and this wasn't exactly a virus. So disabling unnecessary ports is the only other option that makes any sense. Although this should have been done before this point. What should happen is that you remove the affected systems from the network but do NOT shut them off.

<https://www.ftc.gov/business-guidance/resources/data-breach-response-guide-business>


upvoted 2 times

A technician is building a lab to learn storage redundancy techniques. Which of the following is the MOST cost-effective method the technician can deploy?

- A. FCoE
- B. Hardware RAID
- C. Software RAID
- D. JBOD

Suggested Answer: C

Community vote distribution



🗨️ **shanebrown** 9 months, 2 weeks ago

Selected Answer: C

JBOD doesn't provide much redundancy and protection.

upvoted 1 times

🗨️ **surfuganda** 1 year, 3 months ago

Selected Answer: D

D. JBOD (Just a Bunch Of Disks)

JBOD is usually the most cost-effective option because it doesn't involve any specialized hardware or RAID controllers. It simply aggregates individual hard drives without any redundancy or striping. While it lacks redundancy and fault tolerance features, it can still be a suitable option for learning storage concepts and experimenting with redundancy techniques at a lower cost compared to the other options listed.

FCoE (Fibre Channel over Ethernet), Hardware RAID, and Software RAID all involve additional hardware or software components, which can increase the overall cost.

upvoted 1 times

🗨️ **ccoli** 11 months, 2 weeks ago

incorrect, software RAID accomplishes this with only 1 physical disk needed.

upvoted 1 times

🗨️ **gingasaurusrex** 2 years, 2 months ago

Selected Answer: C

C. Software RAID

Software RAID is the most cost-effective method of deploying storage redundancy techniques in a lab environment. Software RAID uses the server's CPU and memory to manage the disk redundancy, and can be implemented without additional hardware or specialized equipment. Most operating systems have built-in support for software RAID, making it easy to set up and configure.

Hardware RAID is a more expensive option, as it requires specialized RAID controllers and storage devices that support RAID. FCoE (Fibre Channel over Ethernet) is a high-performance network protocol used to connect storage devices to servers, but it requires specialized network hardware and is not a cost-effective option for a lab environment. JBOD (Just a Bunch Of Disks) is a storage configuration where disks are presented to the operating system as individual drives, and it does not provide any redundancy or fault tolerance.

upvoted 2 times

🗨️ **Pongsathorn** 2 years, 7 months ago

Selected Answer: C

Software RAID advantages:

- Less expensive
- Fewer components to manage

Disadvantages:

- Less flexible
- Consumes system and OS resources

Hardware RAID advantages:

- Faster
- More flexible
- Does not consume server resources

Disadvantages:

- More expensive
- More components to manage
- More complex

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 114.



upvoted 2 times

  **Pongsathorn** 2 years, 7 months ago

Selected Answer: C

JBOD is no data redundancy, and regardless of the number of disks in the system, if the data is spanned across the disks, the loss of a single disk means the loss of all data.

upvoted 2 times


  **Timock** 2 years, 9 months ago

Selected Answer: C

Software RAID would be the most cost effective as it will use existing HDs to create the environment. Not to mention we are looking to learn storage redundancy. JBOD is the exact opposite of this.

<https://www.techtarget.com/searchstorage/definition/software-RAID-software-redundant-array-of-independent-disk#>


upvoted 1 times

  **szl0144** 2 years, 9 months ago

Selected Answer: C

C is correct

upvoted 1 times

  **hakumai** 2 years, 9 months ago

I think the correct answer is C. JBODs don't have redundancy.

upvoted 2 times

A Linux administrator created a script that will run at startup. After successfully writing the script, the administrator received the following output when trying to execute the script: hash: ./startUp.sh: Permission denied

Which of the following commands would BEST resolve the error message?

- A. `chmod +w startUp.sh`
- B. `chmod 444 startUp.sh`
- C. `chmod +x startUp.sh`
- D. `chmod 466 startUp.sh`

Suggested Answer: C

Community vote distribution

C (100%)

 **surfuganda** 9 months, 3 weeks ago

Selected Answer: C

The error message "Permission denied" indicates that the script lacks execute permission. In Linux, to execute a script, it must have the execute permission bit set.

The correct command to resolve this error would be:

C. `chmod +x startUp.sh`

This command adds the execute permission (+x) to the script file "startUp.sh", allowing it to be executed.

upvoted 1 times

 **EngAbood** 11 months, 2 weeks ago

Selected Answer: C

x = excute

r = read

w = write

upvoted 1 times

 **gingasaurusrex** 1 year, 8 months ago

Selected Answer: C

C. `chmod +x startUp.sh`

The error message "hash: ./startUp.sh: Permission denied" indicates that the script cannot be executed because it does not have the executable permission. To resolve the issue, the administrator should use the `chmod` command to add the executable permission to the script. The correct command is "`chmod +x startUp.sh`". This will add the executable permission for the owner of the file, allowing the script to be executed at startup.

The other options do not address the issue of the script not having the executable permission. `chmod +w` would add the write permission, but this is not necessary for a script to run. `chmod 444` would add read-only permission to the owner, group, and others, but again, this is not necessary for a script to run. `chmod 466` would add read and write permission to the owner and group, but this still would not allow the script to be executed.

upvoted 1 times

Due to a recent application migration, a company's current storage solution does not meet the necessary requirements for hosting data without impacting performance when the data is accessed in real time by multiple users. Which of the following is the BEST solution for this issue?

- A. Install local external hard drives for affected users.
- B. Add extra memory to the server where data is stored.
- C. Compress the data to increase available space.
- D. Deploy a new Fibre Channel SAN solution.

Suggested Answer: D

Community vote distribution

D (100%)

  **gingasaurusrex** 8 months, 3 weeks ago



Selected Answer: D

D. Deploy a new Fibre Channel SAN solution.

The BEST solution for hosting data without impacting performance when accessed in real time by multiple users would be to deploy a new storage solution that can handle the requirements. A Fibre Channel SAN solution is a high-speed storage network that provides dedicated connectivity between servers and storage devices. This type of solution is designed for high-performance applications and can provide the necessary performance for multiple users accessing data in real time.

Installing external hard drives for affected users may provide additional storage, but it would not necessarily improve performance, and it would not address the root cause of the issue. Adding extra memory to the server may improve performance, but it would not necessarily be enough to handle the requirements, and it would not address the issue of storage capacity. Compressing the data to increase available space would not necessarily improve performance, and it would add additional overhead for decompression when accessing the data.

upvoted 1 times

  **Timock** 1 year, 3 months ago

Selected Answer: D

Both SAN and NAS systems are network-based storage solutions aimed at providing multiple users 24/7 access to data on-premises and remotely. Here are some differences between the two approaches.

Type of network: NAS is connected to devices using a LAN or Ethernet network, while a SAN runs on high-speed Fibre channel.

upvoted 1 times

Which of the following, if properly configured, would prevent a user from installing an OS on a server? (Choose two.)

- A. Administrator password
- B. Group Policy Object
- C. Root password
- D. SELinux
- E. Bootloader password
- F. BIOS/UEFI password

Suggested Answer: EF

Community vote distribution

EF (100%)

  **gingasaurusrex** 8 months, 3 weeks ago

Selected Answer: EF

- E. Bootloader password
- F. BIOS/UEFI password



If properly configured, a bootloader password or BIOS/UEFI password would prevent a user from installing an OS on a server. These passwords restrict access to the boot process, preventing unauthorized changes to the system.

An administrator password or root password would not necessarily prevent a user from installing an OS, as these passwords are used to control access to the operating system once it has already been installed.

A Group Policy Object is a configuration tool used in Windows environments to manage security settings and other system configurations for groups of users and computers, but it would not prevent a user from installing an OS.

SELinux is a security feature in Linux that provides mandatory access control for processes and files, but it would not prevent a user from installing an OS.

upvoted 2 times

  **Timock** 1 year, 3 months ago

Selected Answer: EF

Requiring a boot password upon execution of the boot loader will prevent an unauthorized user from entering boot parameters or changing the boot partition. This prevents users from weakening security (e.g. turning off SELinux at boot time).

What is a UEFI password? A UEFI, or BIOS, password is a password that must be entered when the machine is powered on or rebooted in order to continue. Without the password the machine cannot be booted at all – even from external media – and no configuration changes to the UEFI or BIOS settings can be made.

https://www.tenable.com/audits/items/CIS_Red_Hat_EL7_STIG_v2.0.0_STIG.audit:1ee8630516da75be0fabf64c11ab70c6

https://askleo.com/how_do_i_remove_a_bios_password/#

upvoted 1 times

A server administrator is completing an OS installation for a new server. The administrator patches the server with the latest vendor-suggested software, configures DHCP, and verifies all network cables are properly connected in the IDF, but there is no network connectivity. Which of the following is the MOST likely reason for the lack of connectivity?

- A. The VLAN is improperly configured.
- B. The DNS configuration is invalid.
- C. The OS version is not compatible with the network switch vendor.
- D. The HIDS is preventing the connection.

Suggested Answer: A

Community vote distribution

A (100%)

🗨️ 👤 **Fakecon** 6 months, 3 weeks ago

Don't know how technician patched server without network connection 😞
upvoted 3 times

🗨️ 👤 **Timock** 2 years, 3 months ago

Selected Answer: A

An intermediate distribution frame (IDF) is a free-standing or wall-mounted rack for managing and interconnecting a telecommunications cable between end-user devices and the main distribution frame (MDF).

DNS has nothing to do with LAN network connectivity. OS version and Switch vendor is not a thing. And the HIDS is detection not prevention. The only choice left is VLAN.

<https://www.techtarget.com/whatis/definition/intermediate-distribution-frame-IDF#>

upvoted 3 times

Which of the following technologies would allow an administrator to build a software RAID on a Windows server?

- A. Logical volume management
- B. Dynamic disk
- C. GPT
- D. UEFI

Suggested Answer: B

Community vote distribution

B (100%)

  **gingasaurusrex** 8 months, 3 weeks ago

Selected Answer: B

The technology that would allow an administrator to build a software RAID on a Windows server is B. Dynamic disk.

Dynamic disks provide software-based disk management that includes support for creating and managing RAID volumes. Logical volume management (A) is a technology used in Linux operating systems for dynamic disk management. GPT (C) is a partitioning scheme used for modern UEFI-based systems, while UEFI (D) is a firmware interface that replaced BIOS in modern systems.

upvoted 1 times

  **Pongsathorn** 1 year, 1 month ago

Selected Answer: B

Once the OS is deployed, additional partitions may be created. Windows Server uses two different methods of organizing storage: Basic Disks and Dynamic Disks. By default, Windows uses the Basic Disk configuration, even though Dynamic Disk configurations are more flexible.

Dynamic disk attributes:

- Manages simple, spanned, striped (RAID 0), mirrored (RAID 1), and striped with parity (RAID 5) volumes.
- Simple and spanned volumes may be expanded to include more storage capacity.
- Mirrored and RAID 5 configurations can be repaired after a failure.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 135.

upvoted 3 times

  **szl0144** 1 year, 3 months ago

Selected Answer: B

I agree with option B

upvoted 2 times

  **hakumai** 1 year, 3 months ago

I think the correct answer is B.

LVM is a technology used in UNIX.

upvoted 4 times

A technician discovers a Linux server cannot mount an NFS share. The following error message occurs after a mount server:/test /mnt/test command: mount point /mnt/test does not exist.

Which of the following should the technician perform FIRST?

- A. Execute ping server.
- B. Execute mkdir -p /mnt/test.
- C. Ensure the server firewall is turned off.
- D. Ensure server: /test is an exported NFS share.

Correct Answer: B


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

Which of the following BEST describes a warm site?

- A. The site has all infrastructure and live data.
- B. The site has all infrastructure and some data.
- C. The site has partially redundant infrastructure and no network connectivity.
- D. The site has partial infrastructure and some data.

Suggested Answer: B

Community vote distribution



gingasaurusrex 8 months, 2 weeks ago

Selected Answer: B

A warm site is a type of disaster recovery site that provides a balance between cost and recovery time objectives. It is an intermediate solution between a hot site and a cold site.

The BEST description of a warm site is B. The site has all infrastructure and some data. A warm site has all the necessary infrastructure, such as power, cooling, and network connectivity, but only a subset of the data required to resume normal operations. This allows for a faster recovery time than a cold site, which has no infrastructure or data, but slower than a hot site, which has all infrastructure and live data.

upvoted 2 times

Obi_Wan_Jacoby 11 months ago

Selected Answer: B

Answer is B. Page 328 and 329 in the (All in one) CompTia Server+ Certification Exam Guide.

It states that equipment (did NOT state "some or partial") is in place and ready to use, but that software and data are needed to get it active.

upvoted 1 times

Obi_Wan_Jacoby 11 months, 1 week ago

Selected Answer: B

Answer is B. Page 329 in the CompTia Server+ Exam Guide

upvoted 1 times

Pongsathorn 1 year, 1 month ago

Selected Answer: D

Hot sites are designed for an immediate takeover of operations in the event of a disaster. They have all the equipment necessary for the business—servers, workstations, network devices, office furniture, power, and Internet connectivity. A hot site is a fully capable location ready at a moment's notice. Because of this, hot sites are costly—they are essentially a mirror of the primary site, with all associated costs. Data is replicated to the hot site regularly to ensure that it has up-to-date content.

Warm sites contain the necessary space for your data center and business offices and some of the required computer and network hardware. For the site to take over from the primary, some equipment needs to be brought in, and some configuration needs to occur. Data has to be migrated to the warm site. This makes the warm site less expensive upfront than a hot site, but it also requires more time to get the business up and running again if there's a disaster.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 248.

upvoted 3 times

Pongsathorn 1 year, 1 month ago

(cont.)

Cold sites contain the necessary office space, without the essential equipment such as workstations, servers, network devices, and office furniture. Data has to be migrated to the cold site. A cold site is basically just the building, and the rest of the equipment must be brought in before the site can take over operations. Power, HVAC, and physical space are all that are provided.

Type Cost Data On-site equipment

Hot High Immediate replication All

Warm Medium Data migrated Some
Cold Low Data migrated None
Cloud Varies Varies None

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 248.
upvoted 1 times

  **Pongsathorn** 1 year, 1 month ago

Type Cost Data On-site equipment
Hot site High Immediate replication All
Warm site Medium Data migrated Some
Cold site Low Data migrated None
Cloud site Varies Varies None

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 248.
upvoted 1 times

  **RSMCT2011** 1 year, 2 months ago

Selected Answer: B

<https://blog.icorps.com/bid/101789/types-of-disaster-recovery-sites>

Cold Computing Sites - the most simplistic type of disaster recovery site. A cold site consists of elements to provide power and networking capability as well as cooling. It does not include other hardware elements such as servers and storage. The use of a cold site is very limiting to a business since before it can be used, backup data along with some additional hardware must be sent to the site and installed. This will impede workflow.

Warm Computing Sites - contain all the elements of a cold site while adding to them additional elements including storage hardware such as tape or disk drives along with both servers and switches. Warm sites are "ready to go" in one sense, but they still need to have data transported to them for use in recovery should a disaster occur.

Hot Computing Sites - a fully functional backup site that already has important data mirrored to it. This is the ideal disaster recovery site but can be challenging to attain.

upvoted 3 times



A technician is trying to reach abc.intranet.com but is unable to do so by name. The technician is able to reach it by IP address, though. Which of the following is MOST likely misconfigured?

- A. The VLAN
- B. The DNS
- C. The subnet mask
- D. The default gateway

Suggested Answer: *B*

Community vote distribution

B (100%)

  **Timock** 9 months, 2 weeks ago

Selected Answer: B



Def. DNS

upvoted 3 times

Which of the following BEST describes the metric that is used to measure server downtime?

- A. MTTR
- B. BIA
- C. RPO
- D. SLA

Suggested Answer: A

  **AzadOB** 10 months, 2 weeks ago

MTTR Mean Time To Recovery

upvoted 2 times

A server administrator has connected a new server to the network. During testing, the administrator discovers the server is not reachable via server name but can be accessed by IP address. Which of the following steps should the server administrator take NEXT? (Choose two.)

- A. Check the default gateway.
- B. Check the route tables.
- C. Check the hosts file.
- D. Check the DNS server.
- E. Run the ping command.
- F. Run the tracert command.

Suggested Answer: CD

Community vote distribution

CD (100%)

  **gingasaurusrex** 8 months, 2 weeks ago

Selected Answer: CD

The server administrator should take the following steps:

D. Check the DNS server: Since the server can be accessed by IP address, but not by name, it suggests a problem with name resolution. Checking the DNS server settings and ensuring that the correct records are in place is a good next step.

C. Check the hosts file: The hosts file is a local file on the server that can contain static IP to name mappings. Checking the hosts file is a good next step to ensure that the correct name resolution is taking place.

The other options may also be helpful in troubleshooting the issue, but they are less likely to be the root cause of the problem. The ping command (E) and tracert command (F) can be used to test network connectivity and identify any routing issues. Checking the default gateway (A) and route tables (B) can help identify any routing issues that may be preventing access to the server.

upvoted 3 times

  **Pongsathorn** 1 year, 1 month ago

Selected Answer: CD

Obviously, the best answers are C and D.



upvoted 1 times

  **nixonbii** 1 year, 2 months ago

Selected Answer: CD

Agree.

upvoted 1 times

  **Timock** 1 year, 3 months ago

Selected Answer: CD

These are some steps for resolving issues with DNS. A gateway is not required for LAN access so this cannot be the answer.

Solution 1: Check your Computer is correct or not for IP Address

Solution 2: Flush and Reset a client resolver cache and Renew DNS client registration

Solution 3: Edit DNS Lookup Host file entry

<https://morgantechspace.com/2014/03/can-ping-by-ip-address-but-not-by-machine-name.html#>


upvoted 2 times

  **szl0144** 1 year, 3 months ago

Selected Answer: CD

C & D are correct

upvoted 2 times

  **hakumai** 1 year, 3 months ago

Isn't the answer C&D?

upvoted 4 times

A datacenter in a remote location lost power. The power has since been restored, but one of the servers has not come back online. After some investigation, the server is found to still be powered off. Which of the following is the BEST method to power on the server remotely?

- A. Crash cart
- B. Out-of-band console
- C. IP KVM
- D. RDP

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **Obi_Wan_Jacoby** 11 months ago

Selected Answer: B

Answer is in-fact B as noted. Look up iDrac as a OOB option for restoring power "<https://infohub.delltechnologies.com/p/idrac9-virtual-power-cycle-remotely-power-cycle-dell-emc-powerededge-servers/>"

upvoted 1 times

🗨️ 👤 **King2** 1 year, 3 months ago

Why not C?

upvoted 1 times

🗨️ 👤 **lordguck** 1 year, 1 month ago

An OutOfBand Console might refer to ILO/Lights-Out-Management which can not be reached by a IP-KVM

upvoted 1 times

An administrator needs to disable root login over SSH. Which of the following files should be edited to complete this task?

- A. /root/.ssh/sshd_config
- B. /etc/ssh/sshd_config
- C. /root/.ssh/ssh_config
- D. /etc/ssh/ssh_config

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **Grumpy_Old_Coot** 7 months, 3 weeks ago

Easy cheat to remember this.. if the etc file has a d on it (sshd), it is the configuration file(s) for a Daemon (or service).
upvoted 2 times

🗨️ 👤 **King2** 1 year, 9 months ago

Selected Answer: B

Answer is correct.

Reference : <https://www.linuxfordevices.com/tutorials/linux/enable-disable-root-login-in-linux>

upvoted 1 times

Users cannot access a new server by name, but the server does respond to a ping request using its IP address. All the user workstations receive their IP information from a DHCP server. Which of the following would be the best step to perform NEXT?

- A. Run the tracert command from a workstation.
- B. Examine the DNS to see if the new server record exists.
- C. Correct the missing DHCP scope.
- D. Update the workstation hosts file.

Suggested Answer: B

Community vote distribution

B (100%)

🗲️ 👤 **RBL23168** 7 months ago

B. Examine the DNS to see if the new server record exists. Obvious answer for me, can't make any sense of other options.
upvoted 1 times

🗲️ 👤 **hasquaati** 9 months, 2 weeks ago

Selected Answer: B

Stupid question. Not enough information to make a choice. DHCP may or may not be giving correct DNS information in scope. DNS entry may or may not be configured. This is open to interpretation. I am choosing B because I am assuming DHCP scope is correct.
upvoted 1 times

🗲️ 👤 **Pongsathorn** 1 year, 7 months ago

Selected Answer: B

A keyword is "All the user workstations receive their IP information from a DHCP server."
upvoted 4 times

🗲️ 👤 **Isaiah44** 11 months, 2 weeks ago


Correcting the DHCP scope?
upvoted 1 times

A technician has several possible solutions to a reported server issue. Which of the following BEST represents how the technician should proceed with troubleshooting?

- A. Determine whether there is a common element in the symptoms causing multiple problems.
- B. Perform a root cause analysis.
- C. Make one change at a time and test.
- D. Document the findings, actions, and outcomes throughout the process.

Suggested Answer: C

Community vote distribution



Timock Highly Voted 1 year, 3 months ago

Selected Answer: C

The question states that there are several possible solutions so step #4 - Establish a theory has been done. Next is to test the theory changing one thing at a time and test for resolution.

Identify the problem.

Research.

Establish a theory of probable cause.

Test the theory.

Establish a plan of action.

Implement or escalate.

Verify functionality.

Document the solution.

<https://www.techtarget.com/whatis/post/CompTIA-troubleshooting-steps-and-solving-common-computer-problems>

upvoted 9 times

Obi_Wan_Jacoby 11 months, 1 week ago

I agree with answer C as well. In the CompTIA Server+ Exam Guide, it has the Server+ Troubleshooting Methodology slightly different than what you have listed. Here it is out of page 277

Identify the problem.

Establish a theory of probable cause

Test the theory.

Establish a plan of action.

Implement a solution or escalate.

Verify functionality.

Perform root cause analysis

Document the solution.

upvoted 2 times

Pongsathorn Highly Voted 1 year, 1 month ago

Selected Answer: A

When troubleshooting, identify any common elements or similar problems that might span multiple servers or network devices. Such common elements might include a new or updated piece of software, a new device driver, or a new configuration.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 16.

upvoted 5 times

gingasaurusrex Most Recent 8 months, 2 weeks ago

Selected Answer: D

D. Document the findings, actions, and outcomes throughout the process

While all of the options are important troubleshooting techniques, documenting the findings, actions, and outcomes throughout the process is the

best way for the technician to proceed with troubleshooting. This will help ensure that the technician can keep track of what has been tried, what has worked, and what has not, allowing them to quickly identify any patterns or commonalities in the symptoms that may be causing multiple problems. It will also make it easier for the technician to share their findings with other members of the team and help prevent similar issues from occurring in the future.

upvoted 1 times

A server administrator needs to replace a 4U rack-mounted device. Which of the following is the BEST way to install the device to ensure proper safety measures?

- A. Lift with the arms while bending the knees.
- B. Use a server lift.
- C. Use a rolling cart.
- D. Lift the device with the help of another administrator.

Suggested Answer: B

Community vote distribution

B (100%)

  **gingasaurusrex** 1 year, 8 months ago

Selected Answer: B

B. Use a server lift.

Using a server lift is the best way to install a 4U rack-mounted device because it is specifically designed to handle heavy and awkwardly shaped equipment safely. It reduces the risk of injury to the server administrator and damage to the device.

Lifting with the arms while bending the knees and lifting the device with the help of another administrator may still put undue strain on the back and other muscles, increasing the risk of injury. Using a rolling cart may not be as stable as a server lift and could lead to the device falling off the cart during transport.

upvoted 1 times

  **TomRW** 2 years, 1 month ago

I have noticed a few times now that the given answer is actually wrong while the community vote is correct. Just how accurate is this site?

upvoted 1 times

  **brenwins** 10 months, 1 week ago

Not at all accurate. The discussion does help though.

upvoted 1 times

  **Pongsathorn** 2 years, 1 month ago

Selected Answer: B

Servers, especially the multi-unit devices, are very heavy. Tall racks require you to lift a server to or over the level of your head. Proper lifting techniques, including team lifting, are recommended. There are also specialized devices designed to lift servers to the appropriate height for rack mounting.

Hydraulic trolley, one type of equipment used for safely lifting servers to the appropriate height in a rack.

It is essential to ensure to secure the rack to the floor. This prevents servers that are pulled forward on rails from overbalancing the rack and causing it to tip over. Floors must also be rated to properly support the weight of the servers, the rack, and any other IT equipment.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 88.

upvoted 1 times

  **dbl29** 2 years, 2 months ago

Selected Answer: B

B is the answer



upvoted 2 times

  **King2** 2 years, 3 months ago

Selected Answer: B

I think B is correct.

upvoted 3 times

  **Timock** 2 years, 3 months ago



Selected Answer: B

A piece of console furniture called a “crash cart” refers to a mobile server room cart equipped with a secured laptop that is used to provide a direct connection to malfunctioning servers and computers for restoration of crashes. - Not the answer.

4U would require at least 4 admin so not the answer....

Server lifts are used to raise equipment to installation level in a frame (cabinet or rack) or to remove equipment from a frame. Server lifts should only be used on level floors strong enough to support lift and load... Using a lift may be slower than maneuvering by hand; but if worker or equipment safety (usually because of weight) is a factor; using a lift is required.

<https://www.idmworks.com/industry-analysis/data-center-server-lift/>
upvoted 3 times

  **szl0144** 2 years, 3 months ago

Selected Answer: B

Do you think B is the option?

upvoted 3 times

A server technician notices a server is very low on disk space. Upon inspecting the disk utilization, the technician discovers server logs are taking up a large amount of space. There is no central log server. Which of the following would help free up disk space?

- A. Log rotation
- B. Log shipping
- C. Log alerting
- D. Log analysis

Suggested Answer: A

Community vote distribution

A (100%)

 **gingasaurusrex** 8 months, 2 weeks ago


Selected Answer: A

A. Log rotation.

Log rotation is a process of automatically moving old log files to an archive location and creating new log files. It helps to prevent the accumulation of old log files that can take up a large amount of disk space. By implementing log rotation, the server technician can ensure that the server logs are managed and maintained in a more efficient way, freeing up disk space.

Log shipping is a process of forwarding log data to a central log server, which is not relevant to this scenario as there is no central log server. Log alerting is a process of generating alerts when specific events occur in the log files, which is not directly relevant to freeing up disk space. Log analysis is the process of analyzing log data to identify trends, patterns, and potential issues, which also does not directly address the issue of low disk space.

upvoted 2 times

 **Timock** 1 year, 3 months ago

In information technology, log rotation is an automated process used in system administration in which log files are compressed, moved (archived), renamed or deleted once they are too old or too big (there can be other metrics that can apply here).

https://en.wikipedia.org/wiki/Log_rotation#

upvoted 2 times



A new Windows server has been introduced to the network. The server administrator is trying to connect remotely to the server but is unsuccessful. Basic troubleshooting shows the server is visible on the network. Which of the following should be configured on the server to be able to connect remotely? (Choose two.)

- A. Remote desktop connection
- B. Remote procedure call
- C. Remote assistance
- D. Port 443
- E. Port 3389
- F. Port 5900

Suggested Answer: AE

Community vote distribution

AE (100%)

  **Kraken84** 8 months, 2 weeks ago

When considering cost-effectiveness and the need to administer and troubleshoot network problems locally on the servers in a server room, the best option is:

B. Crash cart

A crash cart typically consists of a movable trolley equipped with a monitor, keyboard, mouse, and sometimes other tools or devices. It can be easily wheeled up to a server rack and connected directly to a server for local troubleshooting.

upvoted 1 times

  **gingasaurusrex** 1 year, 2 months ago

Selected Answer: AE

- A. Remote desktop connection
- E. Port 3389

To connect remotely to a Windows server, the server must have the Remote Desktop Connection feature enabled and the port 3389 must be open on the server's firewall to allow incoming remote desktop connections. Therefore, the server administrator should configure Remote Desktop Connection on the server and ensure that port 3389 is open on the server's firewall.

Remote procedure call (RPC) and Remote Assistance are not relevant to this scenario. Port 443 and port 5900 are typically used for other services (HTTPS and VNC, respectively) and are not required for remote desktop connections to a Windows server.

upvoted 2 times

A server room with many racks of servers is managed remotely with occasional on-site support. Which of the following would be the MOST cost-effective option to administer and troubleshoot network problems locally on the servers?

- A. Management port
- B. Crash cart
- C. IP KVM
- D. KVM

Suggested Answer: D

Community vote distribution

D (67%)

B (33%)

  **nixonbii** Highly Voted 2 years, 8 months ago

Selected Answer: D

I think this exam has a problem keeping the terms "local" and "remote" in their own lanes. If you are not physically present, then you are remote. If you are even in the same building with the server room, then you are local. The answer should be D.

upvoted 7 times

  **cj207800** Most Recent 5 months ago

Selected Answer: D

I wanted to go with crash cart but after researching it would appear kvm is actually more cost effective

upvoted 1 times


  **SecNoob27639** 9 months, 1 week ago

Selected Answer: D

While everyone is mentioning Local and Cost Effective, I feel like a major factor is being missed. "...administer and troubleshoot network problems..."

While an IP KVM can operate locally (and also local means "at the server" not just "in the same building," you can have a remote server 50 feet from your desk) it's not going to function properly, if at all, when dealing with network issues on-site. KVM. I stress the network issues piece, because 2 different LLLM/AI's gave me the answer of IP KVM, and I don't want someone else to trip over this question because their research tool was flawed.

upvoted 1 times

  **shanebrown** 9 months, 2 weeks ago

Selected Answer: B

B. Crash Cart. Cheapest tool that can be accessed "Locally"

upvoted 1 times

  **shanebrown** 9 months, 2 weeks ago

Update. Did a check with Chat GPT. KVM is actually cheaper than a Crash Cart.

upvoted 1 times

  **ccoli** 11 months, 1 week ago


Keywords are locally and cost effective. Crash cart is cheapest if you are local.

upvoted 1 times

  **ccoli** 11 months, 1 week ago

They really should not have put locally in the question. I'm guessing they want IP KVM because it many racks of servers so you'd need more than one KVM and only one IP KVM.

upvoted 1 times

  **Kraken84** 1 year, 8 months ago

Selected Answer: B

"managed remotely" on the REG / BUT ? asks "Troubleshoot network problems locally on the servers?"

upvoted 1 times

  **gingasaurusrex** 2 years, 2 months ago

Selected Answer: B

the key word here is LOCALLY and MOST COST EFFECTIVE!

B. Crash cart.

A crash cart is the most cost-effective option for administering and troubleshooting network problems locally on the servers in a server room. It is a mobile cart equipped with a monitor, keyboard, and mouse that can be wheeled up to a server for on-site support. It allows the server administrator to interact with the server as if they were physically present, without having to be physically stationed in the server room.


A management port, IP KVM, and KVM are other options for remote server administration, but they are generally more expensive than a crash cart. Management port provides remote access to server management interfaces, but does not allow for physical interaction with the server. IP KVM and KVM provide remote access to servers with physical interaction capabilities, but they require additional hardware and software that can increase costs.

upvoted 3 times

  **papahawaii** 2 years, 5 months ago

Cost effective method = KVM, IP KVM costs more. And it says locally.

upvoted 1 times

  **Pongsathorn** 2 years, 7 months ago

Selected Answer: D

Management port - irrelevant choice.

Crash cart - if you on-site support occasionally, not worth to have crash cart.

IP KVM - this connection type use for remote drive access, remote console access, OOB and but more cost.

KVM - local access, less cost

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 96.

upvoted 2 times


A technician is attempting to reboot a remote physical Linux server. However, attempts to command a shutdown --now result in the loss of the SSH connection.

The server still responds to pings. Which of the following should the technician use to command a remote shutdown?

- A. A virtual serial console
- B. A KVM
- C. An iDRAC
- D. A crash cart

Suggested Answer: A

Community vote distribution



surfuganda 1 year, 3 months ago

Selected Answer: C

C. iDRAC

The point isn't whether they mention Dell. The point is that iDRAC (which is an out-of-band management system) is the only option.

A. A virtual serial console - would require the host OS to be responsive (it isn't)

B. KVM - would be physically located with the server (which is remote), so no help.

C. iDRAC - typically uses a separate network interface and protocols, such as HTTP, HTTPS, SNMP, and IPMI (Intelligent Platform Management Interface), to provide out-of-band management capabilities. This means that even if the server's operating system or network connectivity is unavailable, iDRAC can still be accessed remotely to perform tasks such as power management, system monitoring, and troubleshooting.

D. A crash cart - would be physically located with the server (which is remote), so no help.

upvoted 1 times

[Removed] 10 months, 3 weeks ago

"The server still responds to pings." How do you figure the server isn't responsive?

upvoted 1 times

Kraken84 1 year, 8 months ago

to use an iDrac would assume the user has dell system

upvoted 1 times

MrS 1 year, 12 months ago

Selected Answer: A

The technician should use a virtual serial console to command a remote shutdown. This is because a virtual serial console allows the technician to access the server's console remotely and execute commands as if they were physically present.

upvoted 4 times

billysunshine 2 years, 2 months ago

I thought a crash cart was local only not for remote use?

upvoted 2 times

gingasaurusrex 2 years, 2 months ago

Selected Answer: D

In addition to the hardware components, a crash cart may also include various diagnostic tools, such as a multimeter, network tester, and cable tester, to assist with troubleshooting network issues. The cart may also have extra storage space for spare parts and tools.

Crash carts are especially useful in situations where remote access is not possible, such as during a network outage or when initial server setup is being performed. They also allow for quicker diagnosis and resolution of problems, as the administrator can physically observe and interact with the server.

When setting up a crash cart, it is important to ensure that it is properly grounded and that all components are securely connected. It is also important to ensure that the cart is easily accessible in case of emergency.

Overall, a crash cart can be a valuable tool for network administrators, providing a cost-effective and efficient way to troubleshoot and resolve server issues.

upvoted 1 times

🗨️ 👤 **Obi_Wan_Jacoby** 2 years, 5 months ago

Well, I am thinking maybe answer A here is the best choice. That being said, iDRAC allows you to access a "Virtual Console" to perform this work, so if it mentioned it was a Dell server (and likely not giving virtual console as an option) then I would go with answer C. fyi, the HP side of things use iLO (Integrated Lights Out) which is like Dell's iDrac.

upvoted 1 times

🗨️ 👤 **DingoHunter** 1 year, 8 months ago

Correct. Dell servers run on Windows.

upvoted 1 times

🗨️ 👤 **lordguck** 2 years, 7 months ago

I go for A; iDrac is Dell specific and KVM is not an IP KVM

upvoted 1 times

🗨️ 👤 **Andrewyounan** 2 years, 9 months ago

It's really hard to decide since the question didn't mention Dell server to pick iDRAC, and KVM might mean Kernabase VM not he IP KVM lol :D

upvoted 1 times

An administrator has been asked to verify that all traffic egressing from a company is secured. The administrator confirms all the information that is sent over the network is encrypted. Which of the following describes the type of traffic being encrypted?

- A. Network encapsulation
- B. Off-site data
- C. Secure FTP
- D. Data in transit

Correct Answer: *D*

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

Which of the following BEST describes a disaster recovery site with a target storage array that receives replication traffic and servers that are only powered on in the event of a disaster?

- A. Cloud
- B. Cold
- C. Hot
- D. Warm

Suggested Answer: C

Community vote distribution



cj207800 6 months ago

Selected Answer: D

D. Warm

A warm disaster recovery site aligns most closely with the scenario described for several reasons:

It has a target storage array that receives replication traffic, indicating some level of data synchronization. Warm sites typically have data synchronized daily or weekly.

The servers are only powered on in the event of a disaster. This is a key characteristic of warm sites, where equipment is partially set up but not fully operational until needed.

Warm sites occupy a middle ground between hot and cold sites in terms of readiness and cost.

Key features of a warm site that match the description include:

Partially established network connectivity

Equipment and hardware are set up but not fully operational

Failover occurs within hours or days, as servers need to be powered on

This setup allows for a balance between cost-effectiveness and recovery speed, which is typical of warm disaster recovery sites

upvoted 1 times

kx7tg4xu 9 months, 2 weeks ago

Selected Answer: B

The cold site best describes the situation described in the problem statement. The reasons for this are as follows

The target storage array is receiving replication traffic:

This indicates that data is being replicated from the primary site to the disaster recovery site.

Servers are only powered on in the event of a disaster:

This means that the servers are not operational during normal times, a characteristic of a cold site.

Cold sites are the lowest cost disaster recovery option, but take the longest time to recover. Data is replicated, but it takes longer to get the system up and running.

upvoted 1 times

Obi_Wan_Jacoby 2 years, 5 months ago

Selected Answer: D

Answer is D. Page 329 in the CompTia Server+ Exam Guide. Hot sites are to be on all the time.

upvoted 3 times

Obi_Wan_Jacoby 2 years, 5 months ago

ahh crud, I'm changing my answer to C. With warm sites, it specifically states that software and data are not in place prior to a disaster. Page 328.

upvoted 4 times

Pongsathorn 2 years, 7 months ago

Selected Answer: D

receives replication traffic and servers that are ONLY powered on in the event of a disaster



upvoted 3 times

Pongsathorn 2 years, 7 months ago

Hot sites are designed for an immediate takeover of operations in the event of a disaster. They have all the equipment necessary for the business—servers, workstations, network devices, office furniture, power, and Internet connectivity. A hot site is a fully capable location ready at a moment's notice. Because of this, hot sites are costly—they are essentially a mirror of the primary site, with all associated costs. Data is replicated to the hot site regularly to ensure that it has up-to-date content.

Warm sites contain the necessary space for your data center and business offices and some of the required computer and network hardware. For the site to take over from the primary, some equipment needs to be brought in, and some configuration needs to occur. Data has to be migrated to the warm site. This makes the warm site less expensive upfront than a hot site, but it also requires more time to get the business up and running again if there's a disaster.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 248.
upvoted 2 times

  **Pongsathorn** 2 years, 7 months ago
(cont.)

Cold sites contain the necessary office space, without the essential equipment such as workstations, servers, network devices, and office furniture. Data has to be migrated to the cold site. A cold site is basically just the building, and the rest of the equipment must be brought in before the site can take over operations. Power, HVAC, and physical space are all that are provided.

Type	Cost	Data	On-site equipment
Hot	High	Immediate replication	All
Warm	Medium	Data migrated	Some
Cold	Low	Data migrated	None
Cloud	Varies	Varies	None

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 248.
upvoted 1 times



  **Pongsathorn** 2 years, 7 months ago

Selected Answer: C

Hot Site

A hot site is a leased facility that contains all the resources needed for full operation. This environment includes computers, raised flooring, full utilities, electrical and communications wiring, networking equipment, and uninterruptible power supplies (UPSs). The only resource that must be restored at a hot site is the organization's data, usually only partially. It should only take a few minutes to bring a hot site to full operation. Although a hot site provides the quickest recovery, it is the most expensive to maintain. In addition, it can be administratively hard to manage if the organization requires proprietary hardware or software. A hot site requires the same security controls as the primary facility and full redundancy, including hardware, software, and communication wiring.

upvoted 2 times

  **Timock** 2 years, 9 months ago

Selected Answer: C

The difference between a hot site and a warm site is that while the hot site provides a mirror of the production data-center and its environment(s), a warm site will contain only servers ready for the installation of production environments.

<https://www.seguetech.com/three-stages-disaster-recovery-sites/#>
upvoted 4 times

An administrator is investigating a physical server that will not boot into the OS. The server has three hard drives configured in a RAID 5 array. The server passes POST, but the OS does not load. The administrator verifies the CPU and RAM are both seated correctly and checks the dual power supplies. The administrator then verifies all the BIOS settings are correct and connects a bootable USB drive in the server, and the OS loads correctly. Which of the following is causing the issue?

- A. The page file is too small.
- B. The CPU has failed.
- C. There are multiple failed hard drives.
- D. There are mismatched RAM modules.
- E. RAID 5 requires four drives.

Suggested Answer: C

Community vote distribution

C (100%)

 **gingasaurusrex** 8 months, 2 weeks ago

Selected Answer: C

Based on the information provided, it is likely that the issue is related to the hard drives or the RAID configuration. The fact that the server passes POST and that the OS loads correctly from a bootable USB drive suggests that the CPU, RAM, and BIOS settings are not the cause of the issue.

Since the server has three hard drives configured in a RAID 5 array, it is possible that one or more of the hard drives has failed, which would prevent the OS from loading. RAID 5 requires at least three hard drives to operate, but can tolerate one drive failure without data loss. However, if two or more drives fail, the array will be unable to function.

It is also possible that there is an issue with the RAID controller or the RAID configuration itself, which could be preventing the OS from loading. In any case, further investigation and testing will be needed to determine the exact cause of the issue.

upvoted 3 times

 **Pongsathorn** 1 year, 1 month ago

Selected Answer: C

RAID 5 Disk Striping with Parity

This RAID 5 design provides a good balance between performance, redundancy, and cost-effective capacity. Data is distributed across the stripes on the disks, as with RAID 0, but additional parity information is also written to the disks, which permits data to be recreated if one of the drives in the array fails. RAID 5 is a good general solution for most servers, though the performance of its writes is dependent on the RAID controller used.

RAID 5 requires a minimum of three HDDs.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 114.

upvoted 1 times

 **Pongsathorn** 1 year, 1 month ago

RAID 6 Disk Striping with Double Parity

The RAID 6 design distributes data across a minimum of four HDDs the same way that RAID 0 and RAID 5 do, but it also distributes parity information across two disks. The result is that a RAID 6 array can recover data even with the failure of two HDDs. Reads are quick, like RAID 5, but writes are slower due to the duplication of the parity data. It is a good general solution as long as the performance hit on write tasks is not a problem for your environment.

RAID 6 requires a minimum of four HDDs.

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 114.

upvoted 1 times

 **Obi_Wan_Jacoby** 11 months ago

Well, they are talking about Raid 5 here, not 6. Raid 5 needs 3 drives but can take 1 going down. If 2 drives fail, then that data is gone, thus no OS to boot to. In this case, the array (Raid 5) has at least 2 drives that have failed. therefore, answer is C

upvoted 2 times

Which of the following load-balancing methods would be BEST described as each connection selecting the next path?

- A. Least connection
- B. Source hashed
- C. Most recently used
- D. Round robin

Suggested Answer: D

Community vote distribution

D (100%)

  **gingasaurusrex** 8 months, 2 weeks ago

Selected Answer: D

Which of the following load-balancing methods would be BEST described as each connection selecting the next path?

- A. Least connection
- B. Source hashed
- C. Most recently used
- D. Round robin

upvoted 1 times

  **Pongsathorn** 1 year, 1 month ago

Selected Answer: D

Round robin - Assigns connections in order

Most recently used (MRU) - Attempts to use existing machines for the workload but can spin up new VMs if needed

Weighted scheduling - Used when the server's hardware capabilities vary, allocating more connections to the more powerful servers

The Official CompTIA Server+ Study Guide (Exam SK0-005) page 244.

upvoted 3 times

Which of the following should be configured in pairs on a server to provide network redundancy?

- A. MRU
- B. SCP
- C. DLP
- D. CPU
- E. NIC

Suggested Answer: E

Community vote distribution

E (100%)

🗳️ **brenwins** 10 months, 1 week ago

Agreed. The answer is E
upvoted 1 times

🗳️ **comptiaboy** 1 year, 8 months ago

Selected Answer: E

Definitely E
upvoted 1 times

🗳️ **gingasaurusrex** 1 year, 8 months ago

Selected Answer: E

The network component that should be configured in pairs on a server to provide network redundancy is the NIC (Network Interface Card).

NIC redundancy can be achieved by installing two NICs on a server and configuring them in a team or bond. This allows for the NICs to work in tandem and provide redundancy in case one of the NICs fails or experiences issues. If one NIC fails, the other NIC can continue to handle the network traffic, which helps to ensure that the server remains connected to the network and available.

While other components listed may have redundancy options or configurations, they do not provide network redundancy in pairs on a server. MRU (Maximum Receive Unit) is a networking term that refers to the largest size of a packet that can be received on a network interface. SCP (Secure Copy Protocol) is a protocol used for securely transferring files between systems. DLP (Data Loss Prevention) is a strategy or technology used to prevent the loss or theft of sensitive data. CPU (Central Processing Unit) is a hardware component that processes instructions in a computer.
upvoted 2 times

🗳️ **fionni92** 1 year, 10 months ago

Selected Answer: E

NIC Teaming.
upvoted 2 times

🗳️ **Obi_Wan_Jacoby** 1 year, 11 months ago

Selected Answer: E

Answer is E. It is called NIC Teaming.
upvoted 2 times

🗳️ **Greedy1985** 1 year, 11 months ago

The NIC should be installed in Pairs to Provide Network redundancy. The key word is Network redundancy.
upvoted 3 times

Which of the following BEST describes a security control that requires validating the user's physical characteristics?

- A. OTP
- B. Biometrics
- C. RFID
- D. Security cameras

Correct Answer: *B*

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


An administrator has set up a GPO to disable USB drive access on a company's servers. The GPO is set up as a computer configuration policy and does not use security groups filtering. The administrator is testing to ensure the policy is working correctly. A server in the accounting OU works as expected, but a server in the sales OU is still allowing access to USB drives. The administrator runs the gpupdate /force command and tries again; however, the server still allows access to the USB drive.

Which of the following is the MOST likely cause?

- A. The WMI filter is not configured properly.
- B. Item-level targeting is not configured properly.
- C. The scope is not configured properly.
- D. Loopback processing is not configured properly.

Suggested Answer: C

Community vote distribution



🗨️ **surfuganda** 9 months, 3 weeks ago

Selected Answer: C

C. The scope is not configured properly. The scope of a GPO determines which objects (such as users or computers) it applies to. Since the GPO is working as expected in the accounting OU but not in the sales OU, and there's no mention of security groups filtering being used, it suggests that the GPO may not be linked to the sales OU or the link is not active. This makes the scope of the GPO the most likely cause, as it might not be properly configured to include the sales OU or there could be an issue with how the GPO is linked to that OU.

upvoted 2 times

🗨️ **MrS** 1 year, 5 months ago

Selected Answer: C

The MOST likely cause of the issue is that C. The scope is not configured properly. The GPO may not be linked to the sales OU or there may be another GPO with higher precedence that is overriding the settings. The administrator should check the Group Policy Management console to ensure that the GPO is linked to the sales OU and that there are no conflicting GPOs with higher precedence.

upvoted 4 times

🗨️ **Obi_Wan_Jacoby** 1 year, 11 months ago

Selected Answer: B

Going with B as stated. A nice detailed article on this can be found here "<https://www.policypak.com/resources/pp-blog/gpo-item-level-targeting/>"

upvoted 3 times


A technician recently applied a patch to a working server. After rebooting, the technician notices the server is unable to connect to a nearby database server. The technician validates a connection can be made to the database from another host.

Which of the following is the best NEXT step to restore connectivity?

- A. Enable HIDS.
- B. Change the service account permissions.
- C. Check the firewall rules.
- D. Roll back the applied patch.

Suggested Answer: D

Community vote distribution



🗨️ 👤 **surfuganda** 9 months, 3 weeks ago

Selected Answer: C

C. Check the firewall rules.

Yes, the server worked before the patch, but we have no indication what kind of patch it was. This is a troubleshooting question. So follow the process. You don't just jump straight to remove the patch without investigating possible causes.

It doesn't say "change" the firewall rules, it says "Check" the firewall rules. This is a low effort step that will yield information, even if nothing has changed. Assume host based firewall on the server could have been modified by the patch. What if the firewall was disabled, or enabled by the patch? Antivirus software updates and uninstalls have been known to disable network cards, and similar such anomalies.

upvoted 2 times

🗨️ 👤 **gingasaurusrex** 1 year, 8 months ago

Selected Answer: D

I am going with D as well

upvoted 3 times

🗨️ 👤 **fionni92** 1 year, 10 months ago

Selected Answer: D

Definitely D. Have to bring server back up before troubleshooting.

upvoted 3 times

🗨️ 👤 **Obi_Wan_Jacoby** 1 year, 11 months ago

Going with answer D

upvoted 2 times

🗨️ 👤 **Greedy1985** 1 year, 11 months ago

I think the answer is D because Server had network connectivity before the patch.

upvoted 4 times

Users report they are unable to access an application after a recent third-party patch update. The physical server that is hosting the application keeps crashing on reboot. Although the update was installed directly from the manufacturer's support website as recommended, it has now been recalled and removed from the website, as the update unintentionally installed unauthorized software after a reboot.

Which of the following steps should the administrator perform to restore access to the application while minimizing downtime? (Choose two.)

- A. Uninstall recent updates.
- B. Reimage the server with a different OS.
- C. Run a port scan to verify open ports.
- D. Enable a GPO to uninstall the update.
- E. Scan and remove any malware.
- F. Reformat the server and restore the image from the latest backup.

Suggested Answer: AE

Community vote distribution

AE (100%)

🗳️ 👤 **MrS** 12 months ago

Selected Answer: AE

To restore access to the application while minimizing downtime, the administrator should A. Uninstall recent updates and E. Scan and remove any malware. Uninstalling the recent update that caused the issue will likely resolve the problem with the server crashing on reboot. Scanning and removing any malware will ensure that the unauthorized software installed by the update is removed from the system.

upvoted 3 times

🗳️ 👤 **gingasaurusrex** 1 year, 2 months ago

Selected Answer: AE

- A. Uninstall recent updates.
- E. Scan and remove any malware.

The first step the administrator should take is to uninstall the recent update, as it is causing the server to crash and preventing access to the application. This will help to restore access to the application without any further downtime.

In addition, the administrator should also scan and remove any malware that may have been inadvertently installed as a result of the unauthorized software that was installed by the recent update. This will help to ensure that the server is secure and prevent any future issues.

Options B, C, D, and F are not necessary or recommended for this situation. Reimaging the server with a different OS or reformatting the server and restoring from backup would result in significant downtime and may not address the root cause of the issue. Running a port scan and enabling a GPO to uninstall the update are not necessary for resolving this issue.

upvoted 3 times


An administrator is only able to log on to a server with a local account. The server has been successfully joined to the domain and can ping other servers by IP address.

Which of the following locally defined settings is MOST likely misconfigured?

- A. DHCP
- B. WINS
- C. DNS
- D. TCP

Suggested Answer: C

Community vote distribution



🗲️ 👤 **surfuganda** 9 months, 3 weeks ago

Selected Answer: C

C. DNS (Domain Name System)

DNS translates domain names to IP addresses. For a server to authenticate using domain credentials, it must be able to resolve the domain controller's name to an IP address. If DNS settings are incorrect and the server cannot resolve the domain controller's name, it would not be able to authenticate domain accounts, even though it could still ping other servers by their IP addresses. This makes DNS the most likely candidate for the misconfiguration causing the issue.

upvoted 1 times

🗲️ 👤 **gingasaurusrex** 1 year, 8 months ago

Selected Answer: C

C. DNS

Since the administrator is only able to log on to the server with a local account, it is likely that the server is having difficulty resolving domain names to IP addresses. This suggests that there may be a misconfiguration in the Domain Name System (DNS) settings on the server.

DNS is responsible for translating domain names into IP addresses, which allows computers to locate resources on the network. If the DNS settings on the server are incorrect, the server may not be able to locate domain resources, including domain controllers, which would prevent the administrator from logging on with a domain account.

DHCP (A) is responsible for assigning IP addresses to devices on the network, and WINS (B) is an older technology used to resolve NetBIOS names to IP addresses. Neither of these services are directly related to logging on to a server with a domain account. TCP (D) is a protocol used to establish connections between devices on a network, and is not likely to be the source of the issue in this scenario.

upvoted 3 times

🗲️ 👤 **TPPASS** 1 year, 8 months ago

Selected Answer: D

should D, I think.

<https://learn.microsoft.com/en-us/troubleshoot/windows-server/identity/cant-sign-the-domain>

upvoted 1 times

A server technician downloaded new firmware from the manufacturer's website. The technician then attempted to install the firmware on the server, but the installation failed, stating the file is potentially corrupt.

Which of the following should the technician have checked prior to installing the firmware?

- A. DLP configuration
- B. MBR failure
- C. ECC support
- D. MD5 checksum

Suggested Answer: *D*

Community vote distribution



RBL23168 7 months, 3 weeks ago

Selected Answer: D

D

MD5 checksum is used to verify the integrity of file
upvoted 1 times



Which of the following is a system that scans outgoing email for account numbers, sensitive phrases, and other forms of PII?

- A. SIEM
- B. DLP
- C. HIDS
- D. IPS

Suggested Answer: B

Community vote distribution

B (100%)

  **kx7tg4xu** 9 months, 2 weeks ago

Selected Answer: B

Data Loss Prevention (DLP) monitors data transmitted externally from an organization with the goal of preventing the release of sensitive or Personally Identifiable Information (PII).

The ability to detect "account numbers, sensitive phrases, and other forms of PII" as described in the problem statement is truly a key function of DLP.

upvoted 1 times

Which of the following BEST describes overprovisioning in a virtual server environment?

- A. Committing more virtual resources to virtual machines than there are physical resources present
- B. Installing more physical hardware than is necessary to run the virtual environment to allow for future expansion
- C. Allowing a virtual machine to utilize more resources than are allocated to it based on the server load
- D. Ensuring there are enough physical resources to sustain the complete virtual environment in the event of a host failure

Suggested Answer: B

Community vote distribution

B (71%)

A (29%)

  **Kabbath1986** 2 months, 1 week ago

Selected Answer: A

The correct answer is:

✓ A. Committing more virtual resources to virtual machines than there are physical resources present

Explanation:

Overprovisioning refers to the practice of allocating more virtual resources (such as CPU, RAM, or storage) to virtual machines (VMs) than the physical host actually has available. In a virtualized environment, hypervisors can sometimes allow VMs to be allocated more resources than are physically available, with the assumption that not all VMs will require their full allocation at the same time. However, if too many VMs demand their full resources simultaneously, it can lead to performance issues or even failures.

This practice can work if the overall load is well managed, but it can be risky if resource contention becomes an issue.

upvoted 1 times

  **Stache** 6 months ago



Selected Answer: B

It's (B). Per the CompTIA Server+ Student Guide:

Overprovision

Systems administrators must ensure that the host server has enough physical hardware resources to meet the needs of the VMs. Therefore, the server is equipped with more hardware resources than the anticipated use requires.

upvoted 1 times

  **cj207800** 6 months, 2 weeks ago

Selected Answer: A

The BEST description of overprovisioning in a virtual server environment is:

A. Committing more virtual resources to virtual machines than there are physical resources present

This accurately describes the concept of overprovisioning in virtualization. Overprovisioning refers to allocating more resources to virtual machines collectively than the physical host actually has available

upvoted 1 times

  **SecNoob27639** 9 months, 1 week ago

Selected Answer: A



I chose A, because B won't lead to an actual issue with the virtual servers.

Using a hard drive as an analogy, if I have 1 TB of data to store and a 2TB drive, I will be fine. If I have 1 TB of data to store, but only a 500 GB drive, I will be over capacity.

Over provisioning then in this case, means you have more virtual servers than the physical hardware can support. Or to go back to our analogy, you only have 500GB of capacity, but you are trying to store 1TB of data.

That said, the question is poorly worded.

upvoted 2 times

  **ccoli** 1 year, 1 month ago

poorly worded. if you're overprovisioning the VMs it is A, if you're overprovisioning the host it is B
upvoted 1 times

🗨️ 👤 **ccoli** 11 months, 1 week ago
They want A for the answer btw.
upvoted 2 times

🗨️ 👤 **AzadOB** 1 year, 4 months ago
Selected Answer: B
To provide more than is necessary.
upvoted 1 times

🗨️ 👤 **EngAbood** 1 year, 5 months ago
Selected Answer: B
overprovisioning = To provide more than is necessary.
upvoted 2 times

🗨️ 👤 **a792193** 1 year, 6 months ago
Selected Answer: B
Correct Answer: B
Systems administrators must ensure that the host server has enough physical hardware resources to meet the needs of the VMs. Therefore, the server is equipped with more hardware resources than the anticipated use requires.
upvoted 2 times

🗨️ 👤 **gingasaurusrex** 2 years, 2 months ago
Selected Answer: A
Which of the following BEST describes overprovisioning in a virtual server environment?

- A. Committing more virtual resources to virtual machines than there are physical resources present
 - B. Installing more physical hardware than is necessary to run the virtual environment to allow for future expansion
 - C. Allowing a virtual machine to utilize more resources than are allocated to it based on the server load
 - D. Ensuring there are enough physical resources to sustain the complete virtual environment in the event of a host failure
- upvoted 2 times

An administrator is alerted to a hardware failure in a mission-critical server. The alert states that two drives have failed. The administrator notes the drives are in different RAID 1 arrays, and both are hot-swappable.

Which of the following steps will be the MOST efficient?

- A. Replace one drive, wait for a rebuild, and replace the next drive.
- B. Shut down the server and replace the drives.
- C. Replace both failed drives at the same time.
- D. Replace all the drives in both degraded arrays.

Suggested Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **K1lroy** 1 year, 5 months ago

Selected Answer: C

why wait

upvoted 2 times

🗨️ 👤 **ccoli** 7 months, 1 week ago

don't cry when the controller fails or one/both arrays fail to rebuild. Real world answer is A.

upvoted 3 times

🗨️ 👤 **gingasaurusrex** 1 year, 8 months ago

Selected Answer: C

The most efficient step would be to replace both failed drives at the same time. This will allow for the quickest recovery time since both drives have failed and need to be replaced to restore data redundancy.

Replacing one drive at a time, waiting for a rebuild, and then replacing the next drive will increase the risk of data loss if another drive fails during the rebuild process. Shutting down the server to replace the drives will result in extended downtime. Replacing all the drives in both degraded arrays is not necessary, as only the two failed drives need to be replaced.

upvoted 3 times

🗨️ 👤 **comptiaboy** 1 year, 8 months ago

Selected Answer: C

Different Raid Arrays

upvoted 2 times


A systems administrator is trying to determine why users in the human resources department cannot access an application server. The administrator reviews the application logs but does not see any attempts by the users to access the application.

Which of the following is preventing the users from accessing the application server?

- A. NAT
- B. ICMP
- C. VLAN
- D. NIDS

Suggested Answer: C

Community vote distribution



🗨️ 👤 **RBL23168** 1 year ago

Selected Answer: C

C for me
upvoted 1 times

🗨️ 👤 **Kraken84** 1 year, 2 months ago

If the users in the human resources department cannot access an application server and there are no logs indicating an attempt to access the application, it suggests there might be network segmentation or isolation in place. The most common method to isolate different departments or sections of a network is using VLANs (Virtual Local Area Networks).

Therefore, the correct answer is:

C. VLAN
upvoted 1 times

🗨️ 👤 **Mareo** 1 year, 8 months ago

Selected Answer: C

If the application logs do not show any attempts by users to access the application server, it is likely that some network-related issue is preventing users from accessing the server.

Out of the options given, the most likely culprit is VLAN (Virtual Local Area Network).
upvoted 3 times

🗨️ 👤 **comptiaboy** 1 year, 8 months ago

Selected Answer: C

Can't be ICMP
upvoted 2 times

🗨️ 👤 **gingasaurusrex** 1 year, 8 months ago

Selected Answer: B

B. ICMP

The fact that the administrator is not seeing any attempts by the users to access the application in the application logs suggests that the issue may be related to network connectivity. ICMP (Internet Control Message Protocol) is a network protocol used to test for connectivity between two devices on a network. If ICMP traffic is being blocked by a firewall or other network device, it could prevent the human resources department users from accessing the application server.

NAT (Network Address Translation) (A) is a technique used to remap IP addresses in a network packet, and would not be directly related to the issue of users not being able to access the application server. VLAN (Virtual Local Area Network) (C) is a logical group of network devices that share a common broadcast domain, and would also not be directly related to the issue. NIDS (Network Intrusion Detection System) (D) is a security technology used to detect network attacks and would not be directly related to the issue of users not being able to access the application server.

upvoted 1 times

  **ccoli** 7 months, 1 week ago

they aren't trying to ping it... only a VLAN would explain why their connection attempts aren't logged.

upvoted 1 times

An administrator discovers a Bash script file has the following permissions set in octal notation:


777

Which of the following is the MOST appropriate command to ensure only the root user can modify and execute the script?

- A. `chmod go-rwx`
- B. `chmod u=rwx`
- C. `chmod u+wx`
- D. `chmod g-rwx`

Suggested Answer: B

Community vote distribution



🗲️ 👤 **Stache** 6 months ago

Selected Answer: A

(A) `chmod go-rwx`

Removing the `rwx` permissions for groups and owners (`go`) would make it so ONLY the root user can modify and execute the script.

*running "`chmod u=rwx`" would change nothing and script's permissions would STILL be 777.

upvoted 2 times

🗲️ 👤 **cj207800** 6 months ago

Selected Answer: A

Ignore last

777

Which of the following is the MOST appropriate command to ensure only the root user can modify and execute the script?

- A. `chmod go-rwx`
 - B. `chmod u=rwx`
 - C. `chmod u+wx`
 - D. `chmod g-rwx`
- upvoted 1 times

🗲️ 👤 **cj207800** 6 months ago

Selected Answer: B

`chmod go-rwx`: This command would remove read, write, and execute permissions for group and others but would leave the owner's permissions unchanged. If the owner is not root, this does not achieve the desired outcome.

upvoted 1 times

🗲️ 👤 **cj207800** 6 months ago

Selected Answer: B

B. `chmod u=rwx`

This command sets the permissions for the user (owner) to read, write, and execute (`rwx`) while removing all permissions for the group and others. After executing this command, the permissions would be changed from 777 to 700, meaning only the root user can access and modify the script.

upvoted 1 times

🗲️ 👤 **surfuganda** 1 year, 3 months ago

Selected Answer: A

A. `chmod go-rwx`

"`chmod`" is the command used to change the file's mode bits.

"go-rwx" specifies that both group (g) and others (o) should have read (r), write (w), and execute (x) permissions removed (-).

This command removes the read, write, and execute permissions from both the group and others, leaving only the user (assuming the root is the owner) with the ability to read, write, and execute the script. This makes the script's permissions 700 in octal notation, ensuring that only the root user (assuming root owns the file) can modify and execute the script.

upvoted 3 times

🗨️ 👤 **K1lroy** 1 year, 9 months ago

Selected Answer: A

But we don't have the g=o= in the available options. Using just "u=rwx" is not going to remove the permissions from group and the others therefore we should remove it with the A: chmod go-rwx

upvoted 2 times

🗨️ 👤 **comptiaboy** 2 years, 2 months ago

Selected Answer: B

Agree with Mareo B

upvoted 2 times

🗨️ 👤 **Mareo** 2 years, 2 months ago

Selected Answer: B

The octal notation "777" represents full permissions (read, write, and execute) for the owner, group, and others. To ensure only the root user can modify and execute the script, we need to remove write and execute permissions for group and others while keeping the same permissions for the owner. Therefore, the most appropriate command would be:

B. chmod u=rwx,g=o= script_file

This command sets the owner's permissions to read, write, and execute (u=rwx) while removing all permissions for the group (g=) and others (o=), effectively restricting modification and execution to the root user only.

upvoted 2 times

🗨️ 👤 **ccoli** 11 months, 1 week ago

Nope, it's A which removes read write and execute from group and others.

upvoted 1 times

A security technician generated a public/private key pair on a server. The technician needs to copy the key pair to another server on a different subnet.

Which of the following is the MOST secure method to copy the keys?

- A. HTTP
- B. FTP
- C. SCP
- D. USB

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **surfuganda** 9 months, 3 weeks ago

Selected Answer: C

C: SCP (Secure Copy Protocol).

SCP (Secure Copy Protocol) is a secure method for transferring files between hosts on a network. It uses SSH (Secure Shell) for data transfer and provides encryption and authentication. This ensures that the keys are transferred securely without being intercepted or compromised.

upvoted 2 times

🗳️ 👤 **RBL23168** 1 year ago

Selected Answer: C

C. SCP. SCP uses SSH, HTTP and FTP are unencrypted.

upvoted 1 times

A server administrator purchased a single license key to use for all the new servers that will be imaged this year. Which of the following MOST likely refers to the licensing type that will be used?

- A. Per socket
- B. Open-source
- C. Per concurrent user
- D. Volume

Suggested Answer: D

Community vote distribution

D (100%)

🗨️ 👤 **kx7tg4xu** 9 months, 2 weeks ago

Selected Answer: D

Volume licensing is a method whereby a company or organization purchases software licenses for multiple computers or servers in bulk. Volume purchases are usually less expensive than individual purchases.

The problem statement states that the server administrator purchased "a single license key to use for all the new servers that will be imaged this year." This is a typical volume license use case.

upvoted 1 times

A technician is attempting to log in to a Linux server as root but cannot remember the administrator password. Which of the following is the LEAST destructive method of resetting the administrator password?

- A. Boot using a Linux live CD and mount the hard disk to /mnt. Change to the /mnt/etc directory. Edit the passwd file found in that directory.
- B. Reinstall the OS in overlay mode. Reset the root password from the install GUI screen.
- C. Adjust the GRUB boot parameters to boot into single-user mode. Run passwd from the command prompt.
- D. Boot using a Linux live CD and mount the hard disk to /mnt. SCP the /etc directory from a known accessible server to /mnt/etc.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ **surfuganda** 9 months, 3 weeks ago

Selected Answer: C

A. While editing the passwd file could reset the root password, it's prone to error and could potentially corrupt the system if not done correctly.

B. Reinstalling the OS in overlay mode is unnecessary and involves reinstalling the entire operating system, which is very destructive and time-consuming for just resetting a password.

***C. Booting into single-user mode via GRUB and using the passwd command to reset the password is a standard method for resetting passwords on Linux systems.

D. Using SCP to copy the /etc directory from a known accessible server to the mounted hard disk could overwrite crucial system files and potentially cause system instability or corruption.

upvoted 2 times

🗳️ **Kraken84** 1 year, 2 months ago

Among the options provided, the least destructive method to reset the administrator password without losing data or altering major system configurations is:

C. Adjust the GRUB boot parameters to boot into single-user mode. Run passwd from the command prompt.

Single-user mode allows you to log in to the system as the root user without requiring a password. Once in single-user mode, the passwd command can be used to change the root password.

upvoted 1 times

🗳️ **gingasaurusrex** 1 year, 8 months ago

Selected Answer: C

The least destructive method of resetting the administrator password would be option C, adjusting the GRUB boot parameters to boot into single-user mode and running passwd from the command prompt. This method allows the technician to reset the root password without making any changes to the system's configuration files or reinstalling the operating system.

Option A involves editing the system's passwd file, which could potentially cause issues with user accounts and permissions.

Option B involves reinstalling the OS, which is a very destructive method and would result in data loss and a complete reset of the system.

Option D involves copying system configuration files from another server, which could potentially introduce new problems and is not a recommended method for resetting the root password.

upvoted 1 times

A technician noted the RAID hard drives were functional while troubleshooting a motherboard failure. The technician installed a spare motherboard with similar specifications and used the original components.



Which of the following should the technician do to restore operations with minimal downtime?

- A. Reinstall the OS and programs.
- B. Configure old drives to RAID.
- C. Reconfigure the RAID.
- D. Install from backup.

Suggested Answer: C

Community vote distribution

C (100%)

  **surfuganda** 9 months, 3 weeks ago

Selected Answer: C

C. Reconfigure the RAID
upvoted 1 times

  **RBL23168** 1 year ago

Selected Answer: C

C. Reconfigure the RAID.

When a motherboard is replaced, especially if it's similar in specifications, the RAID configuration may be affected. RAID (Redundant Array of Independent Disks) configurations are often tied to the specific hardware on which they were configured. Reconfiguring the RAID ensures that the array is recognized and functions correctly with the new motherboard, minimizing downtime and allowing the system to operate with the original set of components.

upvoted 1 times

A server administrator is trying to determine the cause of a slowdown on a database server. Upon investigation, the administrator determines the issue is in the storage subsystem.

Which of the following will MOST likely resolve this issue?

- A. Increasing IOPS by implementing flash storage
- B. Implementing deduplication on the storage
- C. Extending capacity by installing a 4TB SATA disk
- D. Reformatting the disk as FAT32

Suggested Answer: A

Community vote distribution

A (100%)

🗨️ 👤 **surfuganda** 9 months, 3 weeks ago

Selected Answer: A

A. Increasing IOPS by implementing flash storage.

upvoted 1 times

🗨️ 👤 **Kraken84** 1 year, 2 months ago

The most likely option to resolve a slowdown in the storage subsystem, especially for a database server, would be:

A. Increasing IOPS (Input/Output Operations Per Second) by implementing flash storage.

Flash storage (like SSDs) significantly improves IOPS compared to traditional spinning disks, which can greatly benefit the performance of database operations.

upvoted 1 times

A server administrator made a change in a server's BIOS in an attempt to fix an issue with the OS not turning on. However, the change did not successfully correct the issue.

Which of the following should the server administrator do NEXT?

- A. Reinstall the server OS in repair mode while maintaining the data.
- B. Flash the BIOS with the most recent version.
- C. Reverse the latest change made to the server and reboot.
- D. Restart the server into safe mode and roll back changes.

Suggested Answer: C

Community vote distribution

C (100%)

 **gingasaurusrex** 8 months, 2 weeks ago

Selected Answer: C

A server administrator made a change in a server's BIOS in an attempt to fix an issue with the OS not turning on. However, the change did not successfully correct the issue.

Which of the following should the server administrator do NEXT?

- A. Reinstall the server OS in repair mode while maintaining the data.
- B. Flash the BIOS with the most recent version.
- C. Reverse the latest change made to the server and reboot.
- D. Restart the server into safe mode and roll back changes.


upvoted 1 times

Which of the following BEST describes a guarantee of the amount of time it will take to restore a downed service?

- A. RTO
- B. SLA
- C. MTBF
- D. MTTR

Suggested Answer: A

Community vote distribution



gingasaurusrex Highly Voted 1 year, 8 months ago

Selected Answer: A

The guarantee of the amount of time it will take to restore a downed service is typically referred to as the "RTO" or "Recovery Time Objective." RTO is the amount of time within which a system, application, or service must be restored after a disruption in order to avoid unacceptable consequences associated with a break in business continuity.

Option A, RTO, is the correct answer.

Option B, SLA (Service Level Agreement), is an agreement between a service provider and its customers that defines the level of service, availability, and other aspects of the service that the provider agrees to deliver.

Option C, MTBF (Mean Time Between Failures), is the average time that elapses between failures of a system, component, or product.

Option D, MTTR (Mean Time To Repair), is the average time it takes to repair a system, component, or product after a failure has occurred.

upvoted 7 times

Stache Most Recent 6 months ago

Selected Answer: B

Going with (B), a Service Level Agreement (SLA) is a set definitive contract of expectations between provider and customer that "guarantees" recovery times in the event of a disaster.

* An RTO is more of an internal target for the service provider.

upvoted 1 times

surfuganda 9 months, 3 weeks ago

Selected Answer: D

D. MTTR (Mean Time To Restore)

A. RTO (Recovery Time Objective) refers to the targeted duration within which a service or system must be restored after a disaster or failure, but it doesn't specifically quantify the actual time it takes to restore the service.

B. SLA (Service Level Agreement) is a contract between a service provider and a customer, specifying the level of service to be provided, but it may include various metrics beyond just restoration time.

C. MTBF (Mean Time Between Failures) represents the average time elapsed between failures of a system or component, which is different from the time taken to restore the service after a failure.

***D. MTTR represents the average time it takes to repair a failed system or component and restore it to normal operation. It's an important metric in understanding the reliability and maintainability of a system.

upvoted 1 times

RBL23168 9 months, 2 weeks ago

MTTR is Mean Time To Repair, so you are wrong. It's A.

upvoted 1 times

An administrator is investigating several unexpected documents and video files that recently appeared in a network share. The administrator checks the properties of the files and sees the author's name on the documents is not a company employee. The administrator questions the other users, but no one knows anything about the files. The administrator then checks the log files and discovers the FTP protocol was used to copy the files to the server.

Which of the following needs to be done to prevent this from happening again?

- A. Implement data loss prevention.
- B. Configure intrusion detection.
- C. Turn on User Account Control.
- D. Disable anonymous access.

Suggested Answer: D

Community vote distribution

D (100%)

 **gingasaurusrex** 8 months, 2 weeks ago

Selected Answer: D

D. Disable anonymous access.

The fact that the FTP protocol was used to copy the files to the server suggests that the network share has anonymous access enabled. Disabling anonymous access would require users to authenticate before accessing the share, which would prevent unauthorized access and file transfers from external sources.

Data loss prevention (A) and intrusion detection (B) are both important security measures, but they may not necessarily prevent unauthorized access to a network share. User Account Control (C) is a feature in Windows that helps prevent unauthorized changes to a computer, but it is not directly related to preventing unauthorized access to a network share. Therefore, the best option in this scenario would be to disable anonymous access.

upvoted 2 times

A technician is monitoring a server and notices there is only one NIC plugged in, but the server has two. The NIC is oversaturated, and the technician would like to increase the available bandwidth.

Which of the following solutions would be the BEST option to increase the speed of this NIC?

- A. Link aggregation
- B. Heartbeat
- C. Most recently used
- D. Active-active

Correct Answer: A

  **Trossard12** 2 months ago

Selected Answer: A

Link aggregation, also known as port trunking or channel bonding, is a networking technique that combines multiple physical Ethernet links into a single logical link. This process increases bandwidth, improves fault tolerance, and provides redundancy in case one link fails. It's commonly used in scenarios like data centers or between switches and servers where high availability and performance are critical.

upvoted 1 times

An administrator is troubleshooting performance issues on a server that was recently upgraded. The administrator met with users/stakeholders and documented recent changes in an effort to determine whether the server is better or worse since the changes.

Which of the following would BEST help answer the server performance question?

- A. Server performance thresholds
- B. A server baseline
- C. A hardware compatibility list
- D. An application service-level agreement

Correct Answer: *B*

  **Trossard12** 2 months ago

Selected Answer: B

A baseline is the ideal or standard configuration for that node. It is the configuration against which you want to judge that node going forward.
upvoted 1 times


A junior administrator needs to configure a single RAID 5 volume out of four 200GB drives attached to the server using the maximum possible capacity. Upon completion, the server reports that all drives were used, and the approximate volume size is 400GB.

Which of the following BEST describes the result of this configuration?

- A. RAID 0 was configured by mistake.
- B. RAID 5 was configured properly.
- C. JBOD was configured by mistake.
- D. RAID 10 was configured by mistake.

Suggested Answer: D

Community vote distribution



surfuganda Highly Voted 1 year, 3 months ago

Selected Answer: D

INCORRECT: A. RAID 0 was configured by mistake.
(RAID 0 would yield 4 drives x 200GB = 800GB)

INCORRECT: B. RAID 5 was configured properly.
(RAID 5 would yield 1 drive for parity overhead and 3 drives x 200GB = 600GB)

INCORRECT: C. JBOD was configured by mistake.
(JBOD [just a bunch of disks] would yield 4 drives x 200GB = 800GB)

CORRECT: D. RAID 10 was configured by mistake.
(RAID 10 would yield 2 mirrored sets of 2 drives x 200GB = 400GB)
upvoted 7 times

billysunshine Highly Voted 2 years, 2 months ago

Seems to be raid10 to me - 2 lots of mirrored disks, totalling 400MB usable!?
upvoted 6 times

Kraken84 Most Recent 1 year, 8 months ago

If a RAID 5 volume was set up using four 200GB drives and the resulting volume size is approximately 600GB, then:

B. RAID 5 was configured properly.

RAID 5 uses one disk's worth of space for parity, so in this setup with four 200GB drives, the usable space would be the equivalent of three drives (3 x 200GB = 600GB). The fourth drive's capacity would be used for parity.
upvoted 1 times

SecNoob27639 9 months, 1 week ago

Except the question specifically states that the capacity of the volume was only 400GB. This means that RAID 5 was NOT configured properly. Of the 3 remaining options, D is the one that makes the most sense, as A and C would be a volume size of 800GB. D mirrors and stripes the drives, meaning an array would have 1/2 the total disk capacity available.
upvoted 1 times

Alizade 1 year, 9 months ago

Selected Answer: B

B. RAID 5 was configured properly.
upvoted 1 times



hasquaati 1 year, 9 months ago

Selected Answer: D

Go to: <https://www.raid-calculator.com/default.aspx> and calculate the usable storage.

The answer is D.

upvoted 2 times

  **K1lroy** 1 year, 11 months ago

Selected Answer: D

4x200G would result in 400G usable if in raid10 (should be 600G usable in raid5 3x data disk and 1 for parity)



upvoted 4 times

  **gingasaurusrex** 2 years, 2 months ago

Selected Answer: B

Looking at this answer again, 3 of the 4 200gb drives would be combined to 200 Gb of usable storage with raid 5 config, proly leaving one drive in raid 0. This is a total of 400 GB of useable space and therefore is B

upvoted 1 times

  **Mareo** 2 years, 2 months ago

Selected Answer: B



The result of configuring a single RAID 5 volume out of four 200GB drives attached to the server using the maximum possible capacity would be a RAID 5 configuration properly set up.

RAID 5 is a data storage virtualization technology that combines multiple physical disk drive components into one or more logical units for the purposes of data redundancy, performance improvement, or both. It uses block-level striping with distributed parity to achieve redundancy while improving performance by spreading data over multiple disks.

In this case, using four 200GB drives in a RAID 5 configuration with block-level striping and distributed parity, the total available capacity would be approximately 600GB (400GB usable), which is the sum of the capacity of three drives minus the capacity of one drive (i.e., $3 \times 200\text{GB} = 600\text{GB} - 200\text{GB} = 400\text{GB}$).

Therefore, the correct answer is B. RAID 5 was configured properly.

upvoted 1 times

  **Fineb** 2 years, 2 months ago

I think the calculation should be the total capacity of all disks minus 1 disk, which is $800 - 200 = 600\text{GB}$

upvoted 1 times

  **gingasaurusrex** 2 years, 2 months ago

Selected Answer: A

The best answer to this question is A. RAID 0 was configured by mistake.

RAID 5 uses one of the drives for parity information, so the usable capacity of the four 200GB drives in a RAID 5 array should be around 600GB, not 400GB. The fact that the server reports using all four drives but with only 400GB of capacity suggests that the four drives were combined using a RAID 0 configuration instead of RAID 5. In a RAID 0 configuration, the capacity of the drives is combined without any redundancy or parity information, resulting in faster performance but no fault tolerance. Therefore, the best answer is that RAID 0 was configured by mistake.

upvoted 1 times

A backup application is copying only changed files each time it runs. During a restore, however, only a single file is used.

Which of the following backup methods does this describe?




- A. Open file
- B. Synthetic full
- C. Full incremental
- D. Full differential

Suggested Answer: B

Community vote distribution

B (83%)

C (17%)

  **comptiaboy**  1 year, 2 months ago

Selected Answer: B

this should be synthetic full

upvoted 5 times

  **Ricky496**  5 months, 4 weeks ago

Selected Answer: D

This ChatGPT's answer:

Correct Answer: D. Full differential

Explanation:

A differential backup method copies all files that have changed since the last full backup. During a restore, only two components are required: the last full backup and the latest differential backup. This ensures that the restore process uses a single file from the differential backup to bring the data up to date.

Why Not the Other Options?

A. Open file:

Refers to the ability of a backup application to back up files that are in use (open) during the backup process. It does not describe a backup method.

B. Synthetic full:

Combines a previous full backup with subsequent incremental backups to create a "synthetic" full backup. While this method involves incremental data, it does not match the description of using only one file during restore.

C. Full incremental:

An incremental backup captures only the files that have changed since the last backup (full or incremental). To restore, you need the last full backup and all incremental backups in sequence, not just a single file.


upvoted 1 times

  **MrS** 12 months ago

Selected Answer: C

The backup method that BEST describes the scenario is Full Incremental. Full incremental backups copy all files the first time and then only copy changed files in subsequent backups. The provided answer is correct.

upvoted 1 times

  **mark9999** 10 months, 3 weeks ago

But you have to restore all the incrementals from the point of the last full backup which is time consuming. Synthetic full builds more recent "full backups" by using the existing full and incrementals, creating one more recent full backup.

upvoted 1 times

A very old PC is running a critical, proprietary application in MS-DOS. Administrators are concerned about the stability of this computer. Installation media has been lost, and the vendor is out of business.

Which of the following would be the BEST course of action to preserve business continuity?

- A. Perform scheduled chkdsk tests.
- B. Purchase matching hardware and clone the disk.
- C. Upgrade the hard disk to SSD.
- D. Perform quarterly backups.

Suggested Answer: B

Community vote distribution

B (67%)

A (33%)

🗳️ 👤 **SecNoob27639** 9 months, 1 week ago

Selected Answer: B

While it may be difficult to purchase matching equipment, that is the "best" option presented here. Performing scheduled chkdsk will only let you know if an issue is happening, not help mitigate it. That said, tomf is correct, in a real world application, the company would want to transition to a more modern option to ensure minimal disruption when that legacy hardware inevitably stops working.

upvoted 2 times

🗳️ 👤 **PaytoPlay** 9 months, 3 weeks ago

B, best answer for me. We have this exact issue at my work now. We have a few old machines around with outdated software that fail on occasion. Usually something involving memory or drives, but we keep like equipment around for that reason.

upvoted 1 times

🗳️ 👤 **tomf021959** 1 year, 3 months ago

Sound to me like they are, well, in a bad situation. Of these options B is the best. But they really need to be transitioning their "critical process" to something more modern, maybe even recode it in house with modern software.

upvoted 3 times

🗳️ 👤 **Kraken84** 1 year, 8 months ago

The method described is the "Differential" backup method. In a differential backup, only the files that have changed since the last full backup are backed up. However, during a restore, you would use the last full backup and the latest differential backup to restore the system.

Therefore, the correct answer is:

D. Full differential.

upvoted 1 times

🗳️ 👤 **hasquaati** 1 year, 9 months ago

Selected Answer: A

Good luck in purchasing very old equipment and then cloning the system ;)

upvoted 1 times

An administrator is troubleshooting a RAID issue in a failed server. The server reported a drive failure, and then it crashed and would no longer turn on. There are two arrays on the failed server: a two-drive RAID 0 set for the OS, and an eight-drive RAID 10 set for data.

Which of the following failure scenarios MOST likely occurred?

- A. A drive failed in the OS array.
- B. A drive failed and then recovered in the data array.
- C. A drive failed in both of the arrays.
- D. A drive failed in the data array.

Suggested Answer: A

Community vote distribution

A (100%)

🗨️ 👤 **surfuganda** 9 months, 3 weeks ago

Selected Answer: A

The RAID 0 (striping without parity) set for the OS is not fault tolerant.

If a drive fails here, the OS cannot recover.

The RAID 10 set for data has Mirroring (RAID 1) and Striping (RAID 0) which is fault tolerant.

If a drive fails here, the data can fully recover.

upvoted 1 times

🗨️ 👤 **gingasaurusrex** 1 year, 8 months ago

Selected Answer: A

crashing and failing to boot could be an OS or BIOS issue but since this is focused on raid array config failures, the answer has to be the OS array failing causing the system to not boot. If it were an issue with the Raid 10 Array, the system would boot and you would not see the data in the OS. We do not have enough to confirm that they both failed as well.

upvoted 1 times


A security manager is concerned that a rogue employee could boot a server from an outside USB drive. Which of the following actions can be taken to reduce this risk? (Choose two.)

- A. Close unneeded ports.
- B. Disable unneeded physical ports.
- C. Set a BIOS password.
- D. Install a SIEM.
- E. Disable unneeded services.
- F. Install a HIDS.

Suggested Answer: BC

Community vote distribution

BC (100%)

 **surfuganda** 9 months, 3 weeks ago

Selected Answer: BC

B. Disable unneeded physical ports. Disabling USB ports or any other physical ports not required for the server's operation can prevent unauthorized access through bootable external devices. This directly addresses the concern by making it impossible to connect and boot from an outside USB drive.

C. Set a BIOS password. By setting a BIOS (or UEFI) password, you ensure that changes to the server's boot order or other critical firmware settings cannot be made without authentication. This prevents an unauthorized user from configuring the server to boot from an external device, such as a USB drive.

upvoted 1 times

 **Kraken84** 1 year, 2 months ago

To reduce the risk of a rogue employee booting a server from an outside USB drive, the following actions should be taken:

B. Disable unneeded physical ports.

By disabling unneeded physical USB ports, it prevents anyone from plugging in an unauthorized USB drive or any other external device.

C. Set a BIOS password.

Setting a BIOS or UEFI password restricts access to the system's firmware settings, ensuring that boot order and other critical settings can't be changed without the correct password.

So, the correct answers are:

B. Disable unneeded physical ports.

C. Set a BIOS password.

upvoted 1 times

 **gingasaurusrex** 1 year, 8 months ago

Selected Answer: BC

The two actions that can be taken to reduce the risk of a rogue employee booting a server from an outside USB drive are:

B. Disable unneeded physical ports: By disabling physical ports such as USB ports, the ability of a rogue employee to boot a server from an outside USB drive can be significantly reduced.

C. Set a BIOS password: Setting a BIOS password can prevent unauthorized access to the BIOS and boot options, which can help prevent a rogue employee from booting the server from an outside USB drive.

Note that while the other options listed (A, D, E, and F) may be useful security measures, they do not directly address the specific risk of a rogue employee booting a server from an outside USB drive.

upvoted 2 times

A server administrator is setting up a new payroll application. Compliance regulations require that all financial systems logs be stored in a central location.

Which of the following should the administrator configure to ensure this requirement is met?

- A. Alerting
- B. Retention
- C. Shipping
- D. Rotation

Suggested Answer: C

Community vote distribution

C (100%)

 **gingasaurusrex** Highly Voted 8 months, 2 weeks ago

Selected Answer: C

The administrator should configure log shipping to ensure compliance with the regulation. Log shipping is the process of automatically transferring log files from one location to another, typically from a remote server or device to a central log storage location. This ensures that all log files are stored in a central location, which is a requirement in this case.

Alerting refers to sending notifications when specific events or conditions occur, but it does not relate to storing logs in a central location.

Retention refers to how long log files are kept before they are deleted, which is important for compliance and audit purposes, but it does not directly address the requirement to store logs in a central location.

Rotation refers to the practice of moving older log files to an archive location or deleting them, typically to free up space on the server, and does not address the requirement to store logs in a central location.

upvoted 6 times

Which of the following concepts is in use when dual power supplies are connected to different power sources?



- A. Fault tolerance
- B. Active-passive
- C. Component redundancy
- D. Heartbeat
- E. Link aggregation

Suggested Answer: A

Community vote distribution

A (50%)

C (50%)

  **cj207800** 5 months, 1 week ago

Selected Answer: A

Fault tolerance is a broader concept that encompasses various strategies to ensure system reliability and availability, including component redundancy. In the case of dual power supplies connected to different power sources, the primary goal is to create a fault-tolerant system that can continue operating even if one power source fails.

Component redundancy is indeed a part of this setup, as there are two power supply units. However, the key aspect here is not just the redundancy of components, but the ability of the system to maintain operation in the face of a power source failure. This ability to continue functioning despite a fault is the essence of fault tolerance

upvoted 1 times

  **Stache** 6 months ago

Selected Answer: C

This question sucks so hard. Going with (C).

Per the CompTIA Server+ Student Guide:

Component redundancy is within the server or network. For example, a single server may have multiple NICs or multiple power supplies.

* With that said, the above was all under the Fault Tolerance section of the lesson, but since it's so clearly described I'm going with "Component Redundancy". Again, what absolute disingenuous phrasing and options from CompTIA.



upvoted 1 times

  **Ronn_Burgandy** 9 months ago

Selected Answer: A

The most correct answer is A. Fault Tolerance would encompass both the different power supplies AND the different power sources. C. would only include the different power supplies.

upvoted 2 times

  **SecNoob27639** 9 months, 1 week ago

Selected Answer: A

Different Power Sources is the key here. Dual Power Supplies connected to the same power source is redundancy for if one power supply fails. Dual Power Supplies connected to DIFFERENT power sources, as detailed in the question, is redundancy for if the power source itself fails. This would ensure that if one power source (like main power from the Grid) fails/goes out, the secondary supply connected to a different power source (like a building generator) would allow continued operation.

upvoted 2 times

  **tomf021959** 1 year, 3 months ago

Selected Answer: C

"FT ensures minimal and core business operations stay online, whereas redundancy is concerned with duplication of hardware and software"

<https://www.liquidweb.com/kb/high-availability-vs-fault-tolerance/>

upvoted 2 times

A technician is decommissioning a server from a production environment. The technician removes the server from the rack but then decides to repurpose the system as a lab server instead of decommissioning it.

Which of the following is the most appropriate NEXT step to recycle and reuse the system drives?

- A. Reinstall the OS.
- B. Wipe the drives.
- C. Degauss the drives.
- D. Update the IP schema.

Correct Answer: *B*

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

A server administrator wants to ensure a storage array can survive the failure of two drives without the loss of data. Which of the following RAID levels should the administrator choose?

- A. 0
- B. 1
- C. 5
- D. 6

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **SecNoob27639** 9 months, 1 week ago

Selected Answer: D

It's D. SurfUganda has a great explanation of why.
upvoted 1 times

🗳️ 👤 **surfuganda** 1 year, 3 months ago

Selected Answer: D

RAID 6 extends RAID 5 by adding another parity block, resulting in two parity blocks per data stripe. This means it can withstand the failure of up to TWO DRIVES within the array without losing data. This is its most significant advantage over most other RAID levels, providing a higher level of data security.

upvoted 2 times

🗳️ 👤 **gingasaurusrex** 2 years, 2 months ago

Selected Answer: D

To ensure a storage array can survive the failure of two drives without losing data, the server administrator should choose RAID 6. RAID 6 provides two levels of parity protection, which means that data can be reconstructed even if two drives fail simultaneously. This is achieved by storing two sets of parity data across all drives in the array, which allows the array to survive multiple drive failures while maintaining data integrity.

RAID 0 does not provide any redundancy and does not protect against drive failures. RAID 1 provides mirroring of data between two drives, but it only protects against the failure of a single drive. RAID 5 provides a single level of parity protection, which allows the array to survive a single drive failure, but not two failures.

upvoted 2 times


After the installation of an additional network card into a server, the server will not boot into the OS. A technician tests the network card in a different server with a different OS and verifies the card functions correctly.

Which of the following should the technician do NEXT to troubleshoot this issue?

- A. Remove the original network card and attempt to boot using only the new network card.
- B. Check that the BIOS is configured to recognize the second network card.
- C. Ensure the server has enough RAM to run a second network card.
- D. Verify the network card is on the HCL for the OS.

Suggested Answer: D

Community vote distribution



comptiaboy Highly Voted 2 years, 2 months ago

Selected Answer: D

Definitely D.

upvoted 8 times

cj207800 Most Recent 5 months, 1 week ago

Selected Answer: B

going with b

upvoted 1 times

SecNoob27639 9 months, 1 week ago

Selected Answer: B

Going with B, because it's an additional (and lower) level of verification than HCL. If BIOS can't see the NIC, or the BIOS isn't configured to use that NIC, then the HCL won't really matter, because the OS isn't going to see it anyway.

upvoted 1 times

kx7tg4xu 9 months, 2 weeks ago

Selected Answer: D

D. Verify the network card is on the HCL for the OS.

Checking the HCL is an important step in identifying hardware and OS compatibility issues.

The inability of the OS to recognize the network card is most likely the cause of the boot failure.

upvoted 1 times

surfuganda 1 year, 3 months ago

Selected Answer: D

D. Verify the network card is on the HCL for the OS.

upvoted 1 times

EngAbood 1 year, 5 months ago

both b & d correct , so what we do nowwwwwwwwwww :(

upvoted 1 times

ccoli 11 months, 1 week ago

It's B because he should have checked HCL prior to swap and it is an additional NIC thus needs to be enabled in BIOS. that's my guess anyway.

upvoted 1 times

Musa007 2 years ago

I think the reason why it's B is because of PXE booting. Network card wasn't detected so the OS wouldn't boot because they had PXE boot was enabled for these computers.

upvoted 1 times

A server room contains ten physical servers that are running applications and a cluster of three dedicated hypervisors. The hypervisors are new and only have 10% utilization. The Chief Financial Officer has asked that the IT department do what it can to cut back on power consumption and maintenance costs in the data center.


Which of the following would address the request with minimal server downtime?

- A. Unplug the power cables from the redundant power supplies, leaving just the minimum required.
- B. Convert the physical servers to the hypervisors and retire the ten servers.
- C. Reimage the physical servers and retire all ten servers after the migration is complete.
- D. Convert the ten servers to power-efficient core editions.

Suggested Answer: B

Community vote distribution

B (100%)

 **gingasaurusrex** Highly Voted 8 months, 2 weeks ago

Selected Answer: B

Option B, which involves converting the physical servers to the hypervisors and retiring the ten servers, would address the request with minimal server downtime and reduce power consumption and maintenance costs. Since the hypervisors are new and have low utilization, there is ample capacity to virtualize the existing physical servers. This would reduce the number of physical servers, resulting in lower power consumption and maintenance costs.

Option A, which involves unplugging power cables, may result in unexpected downtime and would not reduce power consumption or maintenance costs.

Option C, which involves re-imaging the physical servers and retiring them after the migration is complete, would require significant downtime and would not necessarily reduce power consumption or maintenance costs.

Option D, which involves converting the physical servers to power-efficient core editions, may reduce power consumption, but would not necessarily

ad
upvoted 6 times

A company needs to increase the security controls on its servers. An administrator is implementing MFA on all servers using cost-effective techniques.

Which of the following should the administrator use to satisfy the MFA requirement?

- A. Biometrics
- B. Push notifications
- C. Smart cards
- D. Physical tokens

Suggested Answer: B

Community vote distribution

B (100%)

  **gingasaurusrex** 8 months, 2 weeks ago

Selected Answer: B



Option B, push notifications, would be the most cost-effective solution to implement MFA on all servers. With push notifications, the user receives a notification on their mobile device or computer prompting them to approve or deny access to the server. This method does not require any additional hardware or tokens and can be implemented easily.

Option A, biometrics, may provide a high level of security, but it would require additional hardware such as fingerprint scanners or facial recognition cameras, which can be expensive.

Option C, smart cards, require the purchase and distribution of smart cards to all users, which can be costly and time-consuming to manage.

Option D, physical tokens, such as key fobs or USB tokens, are also costly to purchase and distribute to all users. Additionally, if a user loses their token, it may cause a delay in access to the server until a new token can be distributed.

upvoted 3 times

  **comptiaboy** 8 months, 3 weeks ago

Selected Answer: B

agree with B

upvoted 1 times

  **Mareo** 8 months, 3 weeks ago

Selected Answer: B

Option B, push notifications, is likely the most cost-effective technique for implementing MFA on servers. This option involves a user attempting to log in with their username and password, and then receiving a push notification on their mobile device asking them to confirm the login attempt. This provides an additional layer of authentication beyond just a username and password.

upvoted 2 times

An organization implements split encryption keys for sensitive files. Which of the following types of risks does this mitigate?

- A. Hardware failure
- B. Malware
- C. Data corruption
- D. Insider threat

Correct Answer: *D*


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

An administrator is installing a new file server that has four drive bays available. Which of the following RAID types would provide the MOST storage as well as disk redundancy?

- A. RAID 0
- B. RAID 1
- C. RAID 5
- D. RAID 10

Suggested Answer: C

Community vote distribution



🗳️ 👤 **hasquaati** 9 months, 2 weeks ago

Selected Answer: C

Go to: <https://www.raid-calculator.com/default.aspx> and calculate the RAID configuration. You will notice that RAID 5 will have the most storage, but still have redundancy. Question doesn't ask which solution has the most redundancy.

upvoted 3 times

🗳️ 👤 **K1lroy** 11 months, 2 weeks ago

Selected Answer: C

I believe it should be raid5 as the word "MOST" has been used on storage, not redundancy

upvoted 4 times

🗳️ 👤 **gingasaurusrex** 1 year, 2 months ago

Selected Answer: D

RAID 10 would provide the MOST storage as well as disk redundancy out of the given options.

RAID 0 provides no redundancy and only offers increased performance by striping data across multiple drives.

RAID 1 provides disk redundancy but only half the total storage capacity since each disk is a mirror of the other.

RAID 5 provides a good balance of storage and redundancy by striping data across multiple disks with parity data distributed among the disks. However, it requires at least three drives and can only withstand the failure of one disk.

RAID 10, also known as RAID 1+0, provides both disk redundancy and increased storage capacity by mirroring two drives and then striping data across the mirrored pairs. This configuration requires at least four drives and can withstand the failure of one drive in each mirrored pair.

upvoted 1 times

🗳️ 👤 **RBL23168** 6 months, 3 weeks ago

Wrong. RAID 10 only has half of all of disk space available due to mirroring (2 disks). RAID 3 has 3 disks so that't the right answer doesn't mention best redundancy anywhere

upvoted 1 times

🗳️ 👤 **RBL23168** 6 months, 3 weeks ago

Meant to say RAID 5 lol even i'm confsuing myself

upvoted 1 times

🗳️ 👤 **Fineb** 1 year, 2 months ago

The keyword Most Storage, Raid 5 has 3 disk for storage and Raid 10 has just 2. I think the correct answer should be RAID 5

upvoted 4 times

🗳️ 👤 **Mareo** 1 year, 2 months ago

RAID 10, also known as RAID 1+0, combines both mirroring and striping. It requires a minimum of four drives

upvoted 1 times

🗳️ 👤 **Mareo** 1 year, 2 months ago

On the other hand, RAID 10 requires a minimum of four drives and uses a combination of mirroring and striping, where data is mirrored across pairs of drives and then striped across the mirrored pairs. In a four-drive RAID 10 configuration, only half of the drives (two drives) are available for data storage, as the other half is used for mirroring.

so you could be right.

upvoted 2 times

A company's servers are all displaying the wrong time. The server administrator confirms the time source is correct.

Which of the following is MOST likely preventing the servers from obtaining the correct time?

- A. A firewall
- B. An antivirus
- C. A HIDS
- D. User account control

Correct Answer: A

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

A server administrator encounters some issues with the server OS after applying monthly patches. Which of the following troubleshooting steps should the administrator perform?

- A. Implement rollback procedures.
- B. Upgrade the drivers.
- C. Reinstall the OS.
- D. Reboot the server.

Correct Answer: A

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

An administrator notices that a file server is filling up quickly but has many of the same files in different directories. Which of the following storage management techniques would resolve this issue?

- A. Provisioning
- B. Encryption
- C. Deduplication
- D. Partitioning

Suggested Answer: C

Community vote distribution

C (100%)

 **gingasaurusrex** 8 months, 2 weeks ago

Selected Answer: C

The storage management technique that would resolve this issue is deduplication.

Deduplication is a technique that identifies and eliminates duplicate copies of data in a storage system. This reduces the amount of storage capacity needed and can improve backup and restore times. In this scenario, since the file server has many of the same files in different directories, deduplication can identify those duplicate files and store only one copy, freeing up space on the server.

Option A, provisioning, refers to the process of allocating and managing storage resources for a particular application or service.

Option B, encryption, is a technique used to protect data by encoding it in such a way that only authorized users can read it.

Option D, partitioning, is the process of dividing a physical storage device into multiple logical units, each of which can be used independently. This can be useful for organizing data or separating operating system files from user files, but it does not address the issue of duplicate files taking up storage space.

upvoted 3 times


An administrator is configuring a new server for use as a database server. It will have two mirrored drives to hold the operating system, and there will be three drive bays remaining for storage.

Which of the following RAID levels will yield the BEST combination of available space and redundancy?

- A. RAID 0
- B. RAID 1
- C. RAID 5
- D. RAID 10

Suggested Answer: C

Community vote distribution



surfuganda 9 months, 2 weeks ago

Selected Answer: C

C. RAID 5

upvoted 2 times

hasquaati 1 year, 3 months ago

Selected Answer: C

Answer is C.

You can only install RAID 10 on even number of disks:

<https://www.ibm.com/docs/en/power9/0009-ESS?topic=srl-raid-10-1>.

upvoted 2 times

billysunshine 1 year, 8 months ago

You can't set up raid 10 with 3 discs, has to be raid 5 - answer C

upvoted 2 times

comptiaboy 1 year, 8 months ago

Selected Answer: D

Agree D

upvoted 1 times

gingasaurusrex 1 year, 8 months ago

Selected Answer: C

The RAID level that would yield the BEST combination of available space and redundancy in this scenario is RAID 5.

RAID 0 provides no redundancy and would not be a suitable choice for a database server where data protection is crucial.

RAID 1 provides disk redundancy by mirroring two drives but would result in only half of the total storage capacity being available. In this case, the three remaining drive bays would provide less storage than RAID 5.

RAID 10 provides both disk redundancy and increased storage capacity by combining mirroring and striping. However, since there are only three remaining drive bays, it would not be possible to set up RAID 10.

RAID 5 provides a good balance of storage capacity and disk redundancy by striping data across multiple disks with parity data distributed among the disks. This configuration requires at least three drives and can withstand the failure of one disk. In this scenario, RAID 5 would provide the best balance of available space and redundancy for the remaining three drive bays.

upvoted 4 times

Mareo 1 year, 8 months ago

Selected Answer: D

In this scenario, the administrator has two mirrored drives for the operating system, leaving three drive bays for storage. The best RAID level that would yield a combination of available space and redundancy would be RAID 10.

RAID 5 provides a good balance of available space and redundancy, but requires at least three disks. It uses parity data to allow recovery of data in the event of a single disk failure.

upvoted 2 times

Hosting data in different regional locations but not moving it for long periods of time describes:

- A. a cold site.
- B. data at rest.
- C. on-site retention.
- D. off-site storage.

Suggested Answer: B

Community vote distribution

B (80%)

D (20%)

🗳️ **surfuganda** 9 months, 2 weeks ago

Selected Answer: B

A. A cold site - does not typically have data until the cold site becomes activated.

***B. Data at rest. - the scenario states "Hosting data in different regional locations". Hosted data is AVAILABLE for use, even if it is not moved or accessed for long periods of time. By contrast, stored data is typically archival, or part of a retention requirement.

C. on-site retention - data retention strategies are typically stored data, not hosted data.

D. off-site storage - the scenario states "Hosting data", not "storing data".

upvoted 2 times

🗳️ **hasquaati** 1 year, 3 months ago

Selected Answer: D

Data at Rest refers to data that is in a storage. It could be a Hard Disk Drive or an SSD. I can be accessed all the time, or never. Where it is stored really doesn't matter. It could be on your home PC, a server, a remote location or anywhere else.

The answer is D because the data is being stored at a remote location which is what the question is asking "different regional location".

upvoted 1 times

🗳️ **gingasaurusrex** 1 year, 8 months ago

Selected Answer: B

The term that describes hosting data in different regional locations but not moving it for long periods of time is "data at rest".

Data at rest refers to data that is not actively being used or transmitted, but is instead stored in a particular location or on a particular device. This can include data that is stored on hard drives, flash drives, or other storage media, as well as data that is stored in the cloud or on remote servers.

Option A, a cold site, refers to a backup facility that is not operational until a disaster occurs and is typically used for disaster recovery purposes.

Option C, on-site retention, refers to the practice of storing data on site, typically for regulatory or compliance purposes.

Option D, off-site storage, refers to the practice of storing data in a location other than the primary data center, typically for backup or disaster recovery purposes.

upvoted 2 times

A technician learns users are unable to log in to a Linux server with known-working LDAP credentials. The technician logs in to the server with a local account and confirms the system is functional, can communicate over the network, and is configured correctly. However, the server log has entries regarding Kerberos errors.

Which of the following is the MOST likely source of the issue?

- A. A local firewall is blocking authentication requests.
- B. The users have expired passwords.
- C. The system clock is off by more than five minutes.
- D. The server has no access to the LDAP host.

Suggested Answer: C

Community vote distribution

C (100%)

 **gingasaurusrex** 8 months, 2 weeks ago

Selected Answer: C

The most likely source of the issue in this scenario is that the system clock is off by more than five minutes.

Kerberos is a network authentication protocol that relies on a shared secret between the client and server, as well as a trusted third-party server that provides authentication services. One of the requirements for proper functioning of Kerberos is that the clocks of the client, server, and third-party server must be synchronized within a certain tolerance. If the clocks are out of sync by more than five minutes, authentication requests will fail.

Option A, a local firewall blocking authentication requests, is possible, but unlikely given that the server log has entries regarding Kerberos errors.

Option B, users having expired passwords, would not result in Kerberos errors, but would result in authentication failure messages.

Option D, the server having no access to the LDAP host, is possible, but unlikely given that the technician confirmed that the system is able to communicate over the network and is configured correctly.

upvoted 3 times

A senior administrator instructs a technician to run the following script on a Linux server:

```
for i in {1..65536}; do echo $i; telnet localhost $i; done
```

The script mostly returns the following message: Connection refused. However, there are several entries in the console display that look like this:

80

Connected to localhost -

443


Connected to localhost -

Which of the following actions should the technician perform NEXT?

- A. Look for an unauthorized HTTP service on this server.
- B. Look for a virus infection on this server.
- C. Look for an unauthorized Telnet service on this server.
- D. Look for an unauthorized port scanning service on this server.

Suggested Answer: A

Community vote distribution



🗨️ **Ronn_Burgandy** 9 months ago

Selected Answer: A

Don't overthink this question. Ports 80 and 443 are for HTTP and HTTPS. Checking for an http service would be the obvious FIRST choice. Even though it is a telnet connection the ports are the default for http. Telnet can also connect to any port if there is any service running does not have to be a specific Telnet service.

upvoted 4 times

🗨️ **SecNoob27639** 9 months, 1 week ago

Selected Answer: C

Per a couple of LLM's/AI's it's C. Their reasoning makes sense as well.

This is a telnet script. Ports 80 and 443 are HTTP and HTTPS, respectively. Because there was a telnet response on both, that indicates potential for issues. When asked to "defend their answer" so to speak, both LLM's returned similar answers, best summarized as:

"While unauthorized HTTP services should certainly be investigated if suspected, the immediate concern in this context is to check for the presence of an unauthorized Telnet service given the nature of the script and the connections made."

upvoted 1 times

🗨️ **SecNoob27639** 8 months, 2 weeks ago

I'm wrong, because the LLM's are wrong. And also stupid. Don't blindly trust them. Just because it was a telnet script doesn't mean anything. The responses were on HTTP and HTTPS ports. Go with Mr. Burgandy up there, and stay classy.


upvoted 1 times

A VLAN needs to be configured within a virtual environment for a new VM. Which of the following will ensure the VM receives a correct IP address?

- A. A virtual router
- B. A host NIC
- C. AVPN
- D. A virtual switch
- E. A vNIC

Suggested Answer: E

Community vote distribution



🗳️ 👤 **PaytoPlay** 5 months, 1 week ago

Selected Answer: D

Assuming this is the core switch, all other devices would get an IP's from the subnets created on this switch.
upvoted 1 times

🗳️ 👤 **cj207800** 5 months, 3 weeks ago

Selected Answer: D

VLAN Configuration: Virtual switches are the primary components for implementing VLANs in virtual environments. They allow you to create port groups with specific VLAN IDs, which VMs then connect to.

Network Segmentation: The virtual switch ensures that the VM is placed in the correct network segment (VLAN), which is crucial for receiving the appropriate IP address.

upvoted 1 times

🗳️ 👤 **[Removed]** 10 months, 3 weeks ago

Selected Answer: D

A vNic by itself would not ensure getting an IP, it is of course necessary but not an assurance.

The fact that they bring a VLAN makes a switch the right option, as it's on a switch that you setup tagging.

upvoted 2 times

🗳️ 👤 **ccoli** 11 months, 1 week ago

The correct answer is D. A virtual switch.

A virtual switch is a software-based network device that connects the virtual machines (VMs) in a virtual environment and allows them to communicate with each

other and with the physical network. A virtual switch can also create and manage virtual LANs (VLANs), which are logical segments of a network that separate the

traffic of different VMs or groups of VMs. A VLAN needs a DHCP server to assign IP addresses to the VMs that belong to it. A virtual switch can act as a DHCP

relay agent and forward the DHCP requests from the VMs to the DHCP server on the physical network. This way, the VMs can receive correct IP addresses for

their VLANs123

upvoted 2 times

🗳️ 👤 **[Removed]** 10 months, 3 weeks ago

I believe that you are right. A vNic by itself would not ensure getting an IP, it is of course necessary but not an assurance.

The fact that they bring a VLAN makes a switch the right option, as it's on a switch that you setup tagging.

upvoted 1 times

🗳️ 👤 **ccoli** 11 months, 1 week ago


A virtual router is a software-based network device that routes packets between different networks or subnets. A virtual router can also create and manage VLANs,

but it is not necessary for a VM to receive a correct IP address. A virtual router can be used to provide additional security, redundancy, or load

balancing for the

VMs



upvoted 1 times

  **ccoli** 11 months, 1 week ago

A host NIC is a physical network interface card that connects the host machine to the physical network. A host NIC can also support VLAN tagging, which allows the host machine to communicate with different VLANs on the network. However, a host NIC alone cannot ensure that a VM receives a correct IP address for its

VLAN. The host NIC needs to be connected to a virtual switch that can relay the DHCP requests from the VMs to the DHCP server

upvoted 1 times

  **ccoli** 11 months, 1 week ago

A VPN is a virtual private network that creates a secure tunnel between two or more devices over the internet. A VPN can be used to encrypt and protect the data

traffic of the VMs, but it is not related to the configuration of VLANs or IP addresses. A VPN does not affect how a VM receives a correct IP address for its VLAN14

A vNIC is a virtual network interface card that connects a VM to a virtual switch or a virtual router. A vNIC can also support VLAN tagging, which allows the VM to

communicate with different VLANs on the network. However, a vNIC alone cannot ensure that a VM receives a correct IP address for its VLAN. The vNIC needs to

be connected to a virtual switch or a virtual router that can relay the DHCP requests from the VMs to the DHCP server

upvoted 1 times

  **surfuganda** 1 year, 3 months ago

Selected Answer: E

E. A vNIC

upvoted 1 times

  **Musa007** 2 years ago

The answer is correct. It's A because a vRouter just like a normal router can act as a dhcp server and give out IP addresses. vNICs play a crucial role in virtual network segmentation. They enable the creation of multiple virtual networks within the same physical infrastructure, allowing for network isolation, traffic management, and improved security. BUT obviously assigning it to a VLAN this way is not the standard. I think we over thought this one here. I'm open to yalls thoughts

upvoted 2 times

  **gingasaurusrex** 2 years, 2 months ago

Selected Answer: E

E. A vNIC (virtual network interface card) is what connects a VM to a virtual network, and it is responsible for sending and receiving network traffic. When a VLAN needs to be configured within a virtual environment for a new VM, creating a vNIC with the correct VLAN configuration will ensure that the VM receives a correct IP address within that VLAN. The virtual switch is what connects the vNIC to the physical network, and a virtual router may be necessary if the VLAN is to communicate with other networks or the Internet, but neither of these options specifically ensures that the VM receives a correct IP address. AVPN (Amazon Virtual Private Network) is a networking solution provided by Amazon Web Services and is not relevant to this scenario unless the virtual environment is hosted on AWS. A host NIC (network interface card) is a physical network adapter on the host machine and is not directly related to configuring a VLAN for a virtual environment.

upvoted 2 times

  **Mareo** 2 years, 2 months ago

Selected Answer: E

To ensure that the VM receives a correct IP address within a VLAN in a virtual environment, the VM must be connected to the virtual switch using a vNIC (Virtual Network Interface Card) configured with the appropriate VLAN ID. Therefore, option E is the correct answer.

A virtual router (option A) may be required if the VM needs to communicate with other networks, but it does not guarantee that the VM will receive a correct IP address within the VLAN.

Therefore, option E (a vNIC) is the correct answer to ensure that the VM receives a correct IP address within a VLAN in a virtual environment.

upvoted 4 times

A user logs in to a Linux server and attempts to run the following command:

```
sudo emacs /root/file
```

However, the user gets the following message:


User userid is not allowed to execute 'emacs' on this server.

Which of the following would BEST allow the user to find out which commands can be used?

- A. visudo | grep userid
- B. sudo -l -U userid
- C. cat /etc/passwd
- D. userlist | grep userid

Suggested Answer: B

Community vote distribution



tame_rabbit 1 year ago

Selected Answer: B

The correct command to list sudo privileges for a user is `sudo -l` or `sudo -ll` (for a more detailed list).

Therefore, the best answer is:

B. `sudo -l -U userid`

Despite the typo (1 instead of l), this option is closest to the correct command. The proper command would be:

```
sudo -l -U userid
```

This command will list all commands that the specified user (userid) is allowed to run using sudo on that system. It's the most direct and appropriate way to find out which commands a user can use with sudo.

upvoted 2 times

surfuganda 1 year, 3 months ago

Selected Answer: A

A. `visudo | grep userid`

Explanation:

The visudo command is used to edit the sudoers file, which contains the rules for granting sudo privileges.

By piping the output of visudo to grep userid, the user can search for any rules or configurations related to the specified user (userid).

This command would allow the user to identify any specific sudo rules that might be restricting their ability to execute the emacs command as sudo.

Reviewing the sudoers file in this manner provides a direct insight into the permissions and commands available to the user.

upvoted 2 times

tame_rabbit 1 year, 4 months ago

it should be `sudo -l -U userid` (not 1)

The `sudo -l -U userid` command is used to list the sudo privileges of a specific user (userid). Here's what each part of the command does:

sudo: This command allows users to execute commands with the security privileges of another user, typically the superuser (root). It is often used to perform administrative tasks while logged in as a regular user.

-l: The -l option stands for "list" and is used to list the allowed (or forbidden) commands that the user can execute with sudo privileges.

-U userid: The -U option specifies the user whose privileges you want to inspect. In this case, userid is the placeholder for the username whose sudo privileges you want to check.

upvoted 1 times

 **gingasaurusrex** 2 years, 2 months ago

Selected Answer: B

The BEST option to allow the user to find out which commands can be used is option B: `sudo -l -U userid`.

This command allows the user to run a shell with the privileges of the specified user (-U userid) while preserving their own environment (-i). This means that the user will be able to see the commands they are authorized to run, and also allows them to test running those commands as the authorized user.

Option A (`visudo | grep userid`) runs the visudo command with the intention of displaying the sudoers file contents that have configuration options for the specified user, but it will not show the available commands.

Option C (`cat /etc/passwd`) shows the contents of the `/etc/passwd` file, which contains user account information including usernames, user IDs, home directories, and login shells. However, it will not show the commands that the user can run.

Option D (`userlist | grep userid`) does not seem to be a valid command and would not help the user find the available commands they can run.

upvoted 2 times

A server administrator is connecting a new storage array to a server. The administrator has obtained multiple IP addresses for the array.

Which of the following connection types is the server MOST likely using to connect to the array?

- A. eSATA
- B. USB
- C. FC
- D. iSCSI

Suggested Answer: D

Community vote distribution

D (100%)

 **gingasaurusrex** Highly Voted 8 months, 2 weeks ago

Selected Answer: D

Based on the multiple IP addresses obtained for the storage array, the server is most likely using the iSCSI (Internet Small Computer System Interface) connection type to connect to the array. iSCSI is a protocol that allows SCSI commands to be sent over IP networks, and it typically uses IP addresses to identify and connect to storage devices over a network.

Option A (eSATA) and B (USB) are both direct-attached storage (DAS) connection types that use a cable to connect the storage device directly to the server, rather than over a network. Option C (FC) is Fibre Channel, which is a network protocol for connecting storage devices to servers over a high-speed Fibre Channel network, but it typically uses Fibre Channel addresses (WWNs) rather than IP addresses to identify and connect to devices.

Therefore, the most likely connection type in this scenario is iSCSI (Option D).

upvoted 6 times

A newly installed server is accessible to local users, but remote users are unable to connect. Which of the following is MOST likely misconfigured?

- A. The IP address
- B. The default gateway
- C. The VLAN
- D. The subnet mask

Suggested Answer: B

Community vote distribution

B (50%)

C (50%)

🗳️ 👤 **17f9ef0** 1 week, 1 day ago

Selected Answer: C

Key here is local vs virtual
upvoted 1 times

🗳️ 👤 **Kabbath1986** 2 months, 1 week ago

Selected Answer: B

The MOST likely culprit is:

B. The default gateway

Why the default gateway?

Local users on the same subnet can reach the server directly, so its IP address and subnet mask are probably correct.

Remote users (on different networks) need a proper default gateway on the server side (and on their end) to route packets between subnets.

If the default gateway is wrong or missing, the server won't know how to send replies back to remote clients, even though it accepts local connections fine.

upvoted 1 times

🗳️ 👤 **ReOliveira30** 4 months, 1 week ago

local users on the same subnet as the server will still be able to access it without issue because they don't need to route traffic outside of their local network

upvoted 1 times

🗳️ 👤 **Stache** 6 months, 2 weeks ago

Selected Answer: C

If default gateway is wrong local users would most likely have issues.

upvoted 1 times

🗳️ 👤 **Ronn_Burgandy** 9 months ago

Selected Answer: B

Outside users are going to have to go through a layer 3 device to get to the server. The router would need to PAT/NAT to server. It doesn't matter what VLAN the server is on as long as the default gateway is configured correctly at layer 3 for port forwarding to the server.

upvoted 4 times

🗳️ 👤 **SecNoob27639** 9 months, 1 week ago

Selected Answer: C

If the default gateway is wrong, local users would also potentially have issues. If it's just the VLAN that is wrong, local users connected to that VLAN will have access, but no one outside that VLAN can reach the server. So it's more likely that it's a VLAN issue than it is a default gateway issue.

upvoted 3 times

A company's security team has noticed employees seem to be blocking the door in the main data center when they are working on equipment to avoid having to gain access each time. Which of the following should be implemented to force the employees to enter the data center properly?

- A. A security camera
- B. A mantrap
- C. A security guard
- D. A proximity card

Correct Answer: *B*

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

Users are able to connect to the wireless network, but they are unable to access the internet. The network administrator verifies connectivity to all network devices, and there are no ISP outages. The server administrator removes the old address leases from the active leases pool, which allows users to access the internet.

Which of the following is MOST likely causing the internet issue?

- A. The DHCP exclusion needs to be removed.
- B. The DHCP scope is full.
- C. The DHCP scope options are misconfigured.
- D. The DHCP lease times are too short.
- E. The DHCP reservations need to be configured.

Suggested Answer: B

Community vote distribution

B (100%)

 **gingasaurusrex** 8 months, 2 weeks ago

Selected Answer: B

This is most likely a typo. Based on the given information, the most likely cause of the internet issue is that the DHCP scope was full.

DHCP (Dynamic Host Configuration Protocol) is a network protocol used to automatically assign IP addresses and other network configuration settings to devices on a network. A DHCP scope is a range of IP addresses that are available for assignment to devices on the network.

If the DHCP scope is full, then new devices that try to connect to the network will not be able to get an IP address and other necessary network configuration settings from the DHCP server. In this case, removing old address leases from the active leases pool would free up some IP addresses for new devices to use, allowing them to access the internet.

The other options (A, C, D, and E) could also cause DHCP-related issues, but they are less likely to be the cause of this specific problem based on the information given.

upvoted 2 times

Which of the following commands would MOST likely be used to register a new service on a Windows OS?

- A. set-service
- B. net
- C. sc
- D. services.msc

Suggested Answer: C

Community vote distribution

C (100%)

  **Stache** 6 months ago

Selected Answer: C

Going with (C) on the assumption it's a typo meant for "sc.exe" (Service Control Manager)

* "sc" (alias for Set-Content) cannot create new services, "set-service", "services.msc", and "net" only modify existing services and objects.

upvoted 1 times

  **gingasaurusrex** 8 months, 2 weeks ago



Selected Answer: C

The command that would most likely be used to register a new service on a Windows OS is "sc".

"sc" is a command-line utility that is used to manage Windows services. It can be used to create, modify, start, stop, and delete services, among other things. The "create" command in particular is used to create a new service.

The other options listed (A, B, and D) are not used specifically for registering new services. "set-service" is used to modify the settings of an existing service, "net" is used for managing network resources, and "services.msc" is a graphical user interface (GUI) tool for managing services.

upvoted 3 times

  **TPPASS** 8 months, 3 weeks ago

Selected Answer: C

The SC command is used to configure, query, stop, start, delete, and add system services on the Windows command line. Most of the SC command subcommands apply equally to drivers. If you have the appropriate permissions, the SC command can be used to manage services on both the local and remote systems.

upvoted 4 times


An administrator receives an alert stating a S.M.A.R.T. error has been detected. Which of the following should the administrator run FIRST to determine the issue?

- A. A hard drive test
- B. A RAM test
- C. A power supply swap
- D. A firmware update

Suggested Answer: A

Community vote distribution

A (100%)

 **SecNoob27639** 9 months, 1 week ago

Selected Answer: A

S.M.A.R.T. is specifically to monitor the health of hard drives. So A is the only applicable answer.
upvoted 1 times

A storage administrator needs to implement SAN-based shared storage that can transmit at 16Gb over an optical connection. Which of the following connectivity options would BEST meet this requirement?

- A. Fibre Channel
- B. FCoE
- C. iSCSI
- D. eSATA

Suggested Answer: A

Community vote distribution

A (100%)

 **gingasaurusrex** 8 months, 2 weeks ago

Selected Answer: A

The connectivity option that would best meet this requirement is Fibre Channel.

Fibre Channel is a high-speed network technology commonly used for storage area networks (SANs). It provides high-bandwidth, low-latency connectivity for storage devices over an optical connection. Fibre Channel supports speeds up to 16Gbps and is therefore the best option for meeting the requirement of transmitting at 16Gb over an optical connection.

FCoE (Fibre Channel over Ethernet) is a technology that encapsulates Fibre Channel frames within Ethernet frames for transmission over Ethernet networks, but it requires special hardware and is not widely adopted.

iSCSI (Internet Small Computer System Interface) is a storage networking standard that carries SCSI commands over IP networks. It is commonly used in smaller SAN environments, but it may not provide the same performance and reliability as Fibre Channel.

eSATA (External Serial Advanced Technology Attachment) is an external interface used to connect storage devices to computers, but it is not a networking technology and does not provide the same capabilities as Fibre Channel.

upvoted 4 times


In which of the following media rotation schemes are daily, weekly, and monthly backup media utilized in a first-in, first-out method?

- A. Waterfall
- B. Synthetic full
- C. Tower of Hanoi
- D. Grandfather-father-son

Suggested Answer: D

Community vote distribution



 **Alizade** 9 months, 1 week ago

Selected Answer: D


D. Grandfather-father-son
upvoted 1 times

A company has a data center that is located at its headquarters, and it has a warm site that is located 20mi (32km) away, which serves as a DR location. Which of the following should the company design and implement to ensure its DR site is adequate?

- A. Set up the warm site as a DR cold site.
- B. Set up a DR site that is in the cloud and in the same region.
- C. Set up the warm site as a DR hot site.
- D. Set up a DR site that is geographically located in another region.

Suggested Answer: D

Community vote distribution



cj207800 5 months, 3 weeks ago

Selected Answer: C

The company has already obtained the site the question is how to make it adequate . I am going to say hot site .
upvoted 1 times

surfuganda 9 months, 2 weeks ago

Selected Answer: D

D. Set up a DR site that is geographically located in another region.
upvoted 2 times

gingasaurusrex 1 year, 8 months ago

Selected Answer: C

A company has a data center that is located at its headquarters, and it has a warm site that is located 20mi (32km) away, which serves as a DR location. Which of the following should the company design and implement to ensure its DR site is adequate?

- A. Set up the warm site as a DR cold site.
 - B. Set up a DR site that is in the cloud and in the same region.
 - C. Set up the warm site as a DR hot site.
 - D. Set up a DR site that is geographically located in another region.
- upvoted 1 times

Alizade 1 year, 8 months ago

Selected Answer: D

D. Set up a DR site that is geographically located in another region.

To ensure the company's disaster recovery (DR) site is adequate, it is advisable to set up a DR site that is geographically located in another region. Having a DR site further away from the primary data center reduces the risk of both sites being affected by the same natural disaster or regional issues. This approach increases the likelihood of maintaining business continuity and ensures that critical systems and data can be recovered in case of a disaster at the main site.

upvoted 4 times

Mareo 1 year, 8 months ago

Selected Answer: C

Option C, setting up the warm site as a DR hot site, would be the most appropriate choice to ensure the DR site is adequate.

A hot site is a DR site that has up-to-date and fully operational hardware, software, and data. In the event of a disaster, the hot site can take over operations almost immediately, ensuring minimal downtime and data loss.

Option D, setting up a DR site that is geographically located in another region, would provide additional protection against disasters that affect the primary site and the warm site. However, it would require additional resources and infrastructure to ensure data synchronization and timely failover.

Therefore, setting up the warm site as a DR hot site would provide the quickest and most effective failover in the event of a disaster, while also being a cost-effective solution.