

# CompTIA

## Exam Questions SK0-005

CompTIA Server+ Certification Exam



**NEW QUESTION 1**

An administrator restores several database files without error while participating in a mock disaster recovery exercise. Later, the administrator reports that the restored databases are corrupt and cannot be used. Which of the following would best describe what caused this issue?

- A. The databases were not backed up to be application consistent.
- B. The databases were asynchronously replicated
- C. The databases were mirrored
- D. The database files were locked during the restoration process.

**Answer:** A

**Explanation:**

Application consistent backup is a method of backing up data that ensures the integrity and consistency of the application state. It involves notifying the application to flush its data from memory to disk and quiescing any write operations before taking a snapshot of the data. If the databases were not backed up to be application consistent, they might contain incomplete or corrupted data that cannot be restored properly. References: CompTIA Server+ Certification Exam Objectives<sup>1</sup>, page 12 What is Application Consistent Backup and How to Achieve It<sup>2</sup> Application-Consistent Backups<sup>3</sup>

**NEW QUESTION 2**

Joe, a user in the IT department cannot save changes to a sensitive file on a Linux server. An ls -l shows the following listing;

```
-rw-r--r 1 Ann IT 6780 12 June 2019 filename
```

Which of the following commands would BEST enable the server technician to allow Joe to have access without granting excessive access to others?

- A. chmod 777 filename
- B. chown Joe filename
- C. Chmod g+w filename
- D. chgrp IT filename

**Answer:** C

**Explanation:**

The chmod command is used to change the permissions of files and directories. The g+w option means to grant write permission to the group owner of the file. Since Joe is a member of the IT group, which is also the group owner of the file, this command will allow him to save changes to the file without affecting the permissions of other users. Verified References: [Linux chmod command]

**NEW QUESTION 3**

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

**Answer:** B

**Explanation:**

Log rotation is a process that periodically renames, compresses, and deletes old log files to free up disk space and keep log files manageable. Log rotation can be configured using tools such as logrotate or cron on Linux systems, or using Windows Task Scheduler or PowerShell scripts on Windows systems. Log rotation can also help with log analysis and troubleshooting by making it easier to find relevant information in smaller and more recent log files. References: <https://www.mezmo.com/learn-log-management/what-is-log-rotation-how-does-it-work/> <https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/logman>

**NEW QUESTION 4**

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

**Answer:** B

**Explanation:**

A core installation is a type of installation method that is best suited to meet the company policy that states no one is to log in directly to the server. A core installation is a minimal installation option that is available when deploying some editions of Windows Server. A core installation includes most but not all server roles and features, but does not include a graphical user interface (GUI). A core installation can only be managed remotely using command-line tools such as PowerShell or Windows Admin Center, or using graphical tools such as Server Manager or Remote Desktop from another computer. This reduces the attack surface, resource consumption, and maintenance requirements of the server. A GUI installation is a type of installation method that includes a graphical user interface (GUI) and allows local or remote management using graphical tools or command-line tools. A virtualized installation is a type of installation method that involves creating and running one or more virtual machines on a physical host using a hypervisor such as Hyper-V or VMware. A clone installation is a type of installation method that involves creating an exact copy of an existing server's configuration and data on another server using tools such as Sysprep or Clonezilla. References: <https://www.howtogeek.com/67469/the-beginners-guide-to-shell-scripting-the-basics/> <https://www.howtogeek.com/443611/how-to-encrypt-your-macs-system-drive-removable-devices-and-individual-files/>

<https://www.howtogeek.com/199068/how-to-upgrade-your-existing-hard-drive-in-under-an-hour/>

#### NEW QUESTION 5

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

**Answer:** C

#### Explanation:

The technician should configure RAID 6 for this drive array to meet the server specifications. RAID 6 is a type of RAID level that provides fault tolerance and performance enhancement by using striping and dual parity. Striping means dividing data into blocks and distributing them across multiple disks to increase speed and capacity. Parity means calculating and storing extra information that can be used to reconstruct data in case of disk failure. RAID 6 uses two sets of parity information for each stripe, which are stored on different disks. This way, RAID 6 can handle up to two disk failures without losing any data or functionality. RAID 6 also allows for the least amount of drive space lost to RAID overhead compared to other RAID levels that can handle two disk failures, such as RAID 1+0 or RAID 0+1.

Reference:

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

#### NEW QUESTION 6

The management team has mandated the use of data-at-rest encryption for all data. Which of the following forms of encryption best achieves this goal?

- A. Drive
- B. Database
- C. Folder
- D. File

**Answer:** A

#### Explanation:

Drive encryption is a form of data-at-rest encryption that encrypts the entire hard drive or solid state drive. This means that all the data on the drive, including the operating system, applications, and files, are protected from unauthorized access. Drive encryption is usually implemented at the hardware or firmware level, and requires a password, PIN, or biometric authentication to unlock the drive. Drive encryption is the most comprehensive and secure way to achieve data-at-rest encryption, as it prevents anyone from accessing the data without the proper credentials, even if they physically remove the drive from the server.

References: CompTIA Server+ Study Guide, Chapter 9: Security, page 367.

#### NEW QUESTION 7

After rack mounting a server, a technician must install four network cables and two power cables for the server. Which of the following is the MOST appropriate way to complete this task?

- A. Wire the four network cables and the two power cables through the cable management arm using appropriate-length cables.
- B. Run the four network cables up the left side of the rack to the top of the rack switch.
- C. Run the two power cables down the right side of the rack toward the UPS.
- D. Use the longest cables possible to allow for adjustment of the server rail within the rack.
- E. Install an Ethernet patch panel and a PDU to accommodate the network and power cables.

**Answer:** B

#### Explanation:

This is the most appropriate way to complete the task because it follows the best practices of cable management. Cable management is a process of organizing and securing cables in a rack or a server room to improve airflow, accessibility, safety, and aesthetics. Running the network cables up the left side and the power cables down the right side of the rack helps to avoid cable clutter, interference, and confusion. It also makes it easier to trace and troubleshoot cables if needed. Using appropriate-length cables also helps to reduce cable slack and excess. Wiring the cables through the cable management arm may cause stress and damage to the cables when moving the server in or out of the rack. Using the longest cables possible may create cable loops and tangles that can block airflow and increase fire hazards. Installing an Ethernet patch panel and a PDU (Power Distribution Unit) may be useful for accommodating more network and power cables, but not necessary for a single server. References: <https://www.howtogeek.com/303282/how-to-manage-your-pcs-fans-for-optimal-airflow-and-cooling/>  
<https://www.howtogeek.com/303290/how-to-properly-manage-your-cables/>

#### NEW QUESTION 8

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

**Answer:** A

#### Explanation:

Retention policies are the guidelines that dictate when to delete data backups based on operational or compliance needs. They specify how long, how, where, and in what format the data backups are stored, and who has authority over them. The other options are not directly related to the deletion of data backups.

<https://backup.ninja/news/Database-Backups-101-Backup-Retention-Policy-Considerations>

#### NEW QUESTION 9

A server administrator is currently working on an incident. Which of the following steps should the administrator perform before resolving the issue?

- A. Inform the impacted users.
- B. Make the changes to the system.
- C. Determine the probable causes.
- D. Identify changes to the server.

**Answer: C**

**Explanation:**

The step that the server administrator should perform before resolving the issue is to determine the probable causes. This step is part of the troubleshooting process that follows a logical and systematic approach to identify and solve problems with servers and applications. The troubleshooting process consists of several steps, such as:

- ? Identify the problem: Gather information from various sources, such as users, logs, or alerts, to understand the symptoms and scope of the problem.
- ? Establish a theory of probable cause: Analyze the information and formulate one or more possible causes of the problem based on evidence or experience.
- ? Test the theory to determine cause: Perform tests or experiments to verify or eliminate each possible cause until the root cause is found.
- ? Establish a plan of action to resolve the problem and implement the solution: Design and execute a plan to fix the problem using appropriate tools and techniques.
- ? Verify full system functionality and implement preventive measures: Confirm that the problem is resolved and that no other issues arise as a result of the solution. Implement preventive measures to avoid recurrence of the problem or improve performance.
- ? Document findings, actions, and outcomes: Record the details of the problem, its cause, its solution, and its outcome for future reference or knowledge sharing. References: [CompTIA Server+ Certification Exam Objectives], Domain 6.0: Troubleshooting, Objective 6.1: Given a scenario involving server hardware issues (e.g., power supply failure), troubleshoot using appropriate tools.

**NEW QUESTION 10**

Which of the following are measures that should be taken when a data breach occurs? (Select 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.

**Answer: BE**

**Explanation:**

These are two measures that should be taken when a data breach occurs. A data breach is an unauthorized or illegal access to confidential or sensitive data by an internal or external actor. A data breach can result in financial losses, reputational damage, legal liabilities, and regulatory penalties for the affected organization. Disclosing the incident is a measure that involves informing the relevant stakeholders, such as customers, employees, partners, regulators, and law enforcement, about the nature, scope, and impact of the data breach. Disclosing the incident can help to mitigate the negative consequences of the data breach, comply with legal obligations, and restore trust and confidence. Identifying the exploited vulnerability is a measure that involves investigating and analyzing the root cause and source of the data breach. Identifying the exploited vulnerability can help to prevent further data loss, remediate the security gaps, and improve the security posture of the organization. Restoring the data from backup is a measure that involves recovering the lost or corrupted data from a secondary storage device or location. However, this does not address the underlying issue of how the data breach occurred or prevent future breaches. Disabling unnecessary ports is a measure that involves closing or blocking network communication endpoints that are not required for legitimate purposes. However, this does not address how the data breach occurred or what vulnerability was exploited. Running an antivirus scan is a measure that involves detecting and removing malicious software from a system or network. However, this does not address how the data breach occurred or what vulnerability was exploited. Moving the data to a different location is a measure that involves transferring the data to another storage device or location that may be more secure or less accessible. However, this does not address how the data breach occurred or what vulnerability was exploited. References: <https://www.howtogeek.com/428483/what-is-end-to-end-encryption-and-why-does-it-matter/> <https://www.howtogeek.com/202794/what-is-the-difference-between-127.0.0.1-and-0.0.0.0/> <https://www.howtogeek.com/443611/how-to-encrypt-your-macs-system-drive-removable-devices-and-individual-files/>

**NEW QUESTION 10**

An organization recently experienced power outages. The administrator noticed the server did not have enough time to shut down properly. After the outages, the administrator had additional batteries installed in the UPS. Which of the following best describes the solution the administrator implemented?

- A. The solution reduced shutdown time.
- B. The solution improved load balancing.
- C. The solution increased power out.
- D. The solution extended runtime.

**Answer: D**

**Explanation:**

The solution the administrator implemented extended runtime. Runtime is the amount of time that a UPS can provide backup power to a server in case of a power outage. By installing additional batteries in the UPS, the administrator increased the capacity and duration of the backup power, allowing the server more time to shut down properly. References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 1, Lesson 1.4, Objective 1.4

**NEW QUESTION 13**

An organization is donating its outdated server equipment to a local charity. Which of the following describes what the organization should do BEFORE donating the equipment?

- A. Remove all the data from the server drives using the least destructive method.
- B. Repurpose and recycle any usable server components.
- C. Remove all the components from the server.
- D. Review all company policies.

**Answer: D**

**Explanation:**

Before donating the outdated server equipment to a local charity, the organization should review all company policies regarding data security, asset disposal, and social responsibility. This can help ensure that the donation complies with the legal and ethical standards of the organization and does not pose any risk to its reputation or operations. Verified References: [Data security], [Asset disposal], [Social responsibility]

**NEW QUESTION 18**

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

**Answer: D**

**Explanation:**

A misconfigured NIC on the server is the most likely reason for the lack of network connectivity. The output of the ping command shows that the server is unable to reach its default gateway (10.0.0.1) or any other IP address on the network. The output of the ipconfig command shows that the server has a valid IP address (10.0.0.10) and subnet mask (255.255.255.0) but no default gateway configured. This indicates that there is a problem with the NIC settings on the server, such as an incorrect IP address, subnet mask, default gateway, DNS server, etc. A misconfigured NIC can also cause an amber link light on the switch port, which indicates a speed or duplex mismatch between the NIC and the switch.

**NEW QUESTION 22**

Which of the following environmental controls must be carefully researched so the control itself does not cause the destruction of the server equipment?

- A. Humidity control system
- B. Sensors
- C. Fire suppression
- D. Heating system

**Answer: C**

**Explanation:**

Fire suppression systems are designed to extinguish or contain fires in a server room, but they can also damage the server equipment if they are not carefully researched and selected. For example, water-based fire suppression systems can cause electrical shorts and corrosion, while gas-based fire suppression systems can create thermal shock and reduce oxygen levels. Therefore, fire suppression systems must be compatible with the server environment and equipment. References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 1, Lesson 1.5, Objective 1.5

**NEW QUESTION 24**

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

**Answer: C**

**Explanation:**

The first thing that the technician should ensure to correct the database synchronization errors is that NTP is running on the database system. NTP (Network Time Protocol) is a protocol that synchronizes the clocks of network devices with a reference time source, such as an atomic clock or a GPS receiver. NTP ensures that all devices on a network have accurate and consistent time settings, which can affect various functions and applications. Database synchronization is a process of maintaining data consistency and integrity across multiple database servers or instances. Database synchronization can depend on accurate time settings, as time stamps are often used to determine which data is newer or older, and which data should be updated or overwritten. If NTP is not running on the database system, it may cause time drift or discrepancy between different database servers or instances, which can result in synchronization errors or data conflicts.

**NEW QUESTION 27**

A server administrator receives the following output when trying to ping a local host:

```
ping imhrh-vc.net
Reply from imhrh-vc.net. Destination host unreachable.
Reply from imhrh-vc.net. Destination host unreachable.
Reply from imhrh-vc.net. Destination host unreachable.
Reply from imhrh-vc.net. Destination host unreachable.
```

Which of the following is MOST likely the issue?

- A. Firewall
- B. DHCP
- C. DNS
- D. VLAN

**Answer:** A

**Explanation:**

A firewall is a network device or software that filters and controls the incoming and outgoing traffic based on predefined rules. A firewall can block or allow certain types of packets, ports, protocols, or IP addresses. The output of the ping command shows that the local host is unreachable, which means that there is no network connectivity between the source and the destination. This could be caused by a firewall that is blocking the ICMP (Internet Control Message Protocol) packets that ping uses to test the connectivity. References: <https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-objectives> (Objective 2.2)

**NEW QUESTION 32**

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

**Answer:** D

**Explanation:**

A port scanner is the tool that would most likely help the analyst determine if the applications are running on a remote server. A port scanner is a software tool that scans a network device for open ports. Ports are logical endpoints for network communication that are associated with specific applications or services. By scanning the ports on a remote server, the analyst can identify what applications or services are running on that server and what protocols they are using. A port scanner can also help detect potential vulnerabilities or misconfigurations on a server.

**NEW QUESTION 34**

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

**Answer:** B

**Explanation:**

The next thing that the technician should do, given the disk is hot swappable, is to replace the disk. A hot swappable disk is a disk that can be removed and replaced without shutting down the system or affecting its operation. A hot swappable disk is typically used in a storage array that has RAID (Redundant Array of Independent Disks) configuration that provides fault tolerance and redundancy. If a disk fails in a RAID array, it can be replaced by a new disk without interrupting the service or losing any data. The new disk will automatically rebuild itself using the data from the other disks in the array.

**NEW QUESTION 39**

A company wants to find an affordable way to simulate a fail over of a critical application. The company does not currently have a solution for it. The application consists of 15 servers, and the company would like to simulate on production configurations and IP address schemes. Which of the following would be the most cost-effective solution?

- A. Build a warm site and perform a fail over of the application.
- B. Build a cloud IaaS and perform a fail over of the application.
- C. Build a hot site and perform a fail over of the application.
- D. Build a cold site and perform a fail over of the application.
- E. Perform a tabletop fail over of the application.

**Answer:** B

**Explanation:**

Cloud IaaS (Infrastructure as a Service) is a service model that allows users to rent virtualized computing resources over the internet, such as servers, storage, network, and software. Cloud IaaS can provide several benefits for disaster recovery and failover scenarios, such as:

- ? Lower cost: Cloud IaaS can reduce the capital and operational expenses of building and maintaining a physical disaster recovery site, as users only pay for the resources they use on demand<sup>12</sup>.
- ? Scalability: Cloud IaaS can offer flexible and elastic scalability of resources, as users can easily provision or deprovision resources according to their needs and workload<sup>12</sup>.
- ? Availability: Cloud IaaS can ensure high availability and reliability of the

application, as users can leverage the cloud provider's redundant and geographically distributed infrastructure<sup>12</sup>.

? Simplicity: Cloud IaaS can simplify the failover process, as users can use the cloud provider's tools and services to automate and orchestrate the failover operations<sup>12</sup>.

Therefore, building a cloud IaaS and performing a failover of the application would be the most cost-effective solution for the company, as it would allow them to simulate a failover of a critical application on production configurations and IP address schemes without investing in a physical disaster recovery site.

#### NEW QUESTION 44

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

**Answer:** D

#### Explanation:

Grandfather-father-son (GFS) is a common backup rotation scheme that uses daily, weekly, and monthly backup media in a first-in, first-out (FIFO) method. The daily backups are rotated on a 3-months basis using a FIFO system as above. The weekly backups are similarly rotated on a bi-yearly basis, and the monthly backups are rotated on an annual basis. The oldest backup media in each cycle are overwritten by the newest ones. This scheme provides multiple versions of backup data at different intervals, allowing for flexible restoration options. Waterfall is another name for GFS. Synthetic full is a backup method that combines an initial full backup with subsequent incremental backups to create a new full backup without transferring all data again. Tower of Hanoi is another backup rotation scheme that uses an algorithm based on moving disks between three pegs. References:

? [https://en.wikipedia.org/wiki/Backup\\_rotation\\_scheme](https://en.wikipedia.org/wiki/Backup_rotation_scheme)

#### NEW QUESTION 47

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.

**Answer:** BDE

#### Explanation:

To comply with the user requirements, the best options are to install the OS on a RAID 1 array, configure RAID 5 for the application data, and use SSD hard drives for the data application array. Here is why:

? RAID 1 is a mirroring technique that creates an exact copy of data on two disks.

This provides redundancy and fault tolerance in case of hard drive failure. RAID 1 also improves read performance since either disk can be read at the same time. Therefore, installing the OS on a RAID 1 array meets the first requirement of separating the OS from the application data and protecting it from hard drive failure.

? RAID 5 is a striping technique with parity that distributes data and parity blocks

across three or more disks. This provides improved performance and storage efficiency compared to RAID 1, as well as fault tolerance in case of a single disk failure. Therefore, configuring RAID 5 for the application data meets the second and third requirements of providing high IOPS performance and data availability.

? SSD hard drives are solid-state drives that use flash memory to store data. They

have no moving parts and offer faster read and write speeds, lower latency, and lower power consumption than traditional HDDs. Therefore, using SSD hard drives for the data application array meets the second requirement of providing high IOPS performance.

References:

? <https://phoenixnap.com/kb/raid-levels-and-types>

? [https://en.wikipedia.org/wiki/Standard\\_RAID\\_levels](https://en.wikipedia.org/wiki/Standard_RAID_levels)

#### NEW QUESTION 52

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. 1
- B. 5
- C. 6

**Answer:** D

#### Explanation:

RAID 6 is a level of RAID that can survive the failure of two drives without the loss of data. RAID 6 uses block-level striping with two parity blocks distributed across all member disks. RAID 6 can tolerate two simultaneous drive failures and still provide data access and redundancy. RAID 0 is a level of RAID that uses striping without parity or mirroring, and offers no fault tolerance. RAID 0 cannot survive any drive failure without data loss. RAID 1 is a level of RAID that uses mirroring without parity or striping, and offers fault tolerance by duplicating data on two or more disks. RAID 1 can survive one drive failure without data loss, but not two. RAID 5 is a level of RAID that uses block-level striping with one parity block distributed across all member disks. RAID 5 can tolerate one drive failure without data loss, but not two. References:

? [https://en.wikipedia.org/wiki/Standard\\_RAID\\_levels](https://en.wikipedia.org/wiki/Standard_RAID_levels)

#### NEW QUESTION 54

Which of the following is an example of load balancing?

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

**Answer:** A

**Explanation:**

Round robin is an example of load balancing. Load balancing is the method of distributing network traffic equally across a pool of resources that support an application. Load balancing improves application availability, scalability, security, and performance by preventing any single resource from being overloaded or unavailable. Round robin is a simple load balancing algorithm that assigns each incoming request to the next available resource in a circular order. For example, if there are three servers (A, B, C) in a load balancer pool, round robin will send the first request to server A, the second request to server B, the third request to server C, the fourth request to server A again, and so on. Reference: <https://simplicable.com/new/load-balancing>

**NEW QUESTION 56**

An application needs 10GB of RAID 1 for log files, 20GB of RAID 5 for data files, and 20GB of RAID 5 for the operating system. All disks will be 10GB in capacity. Which of the following is the MINIMUM number of disks needed for this application?

- A. 6
- B. 7
- C. 8
- D. 9

**Answer:** C

**Explanation:**

To calculate the minimum number of disks needed for this application, we need to consider the RAID levels and their disk requirements. RAID 1 requires a minimum of two disks and provides mirroring, which means that data is duplicated on both disks. RAID 5 requires a minimum of three disks and provides striping with parity, which means that data is distributed across all disks with one disk storing parity information for error correction. RAID 5 can tolerate one disk failure without losing data. To create a 10GB RAID 1 array for log files, we need two 10GB disks. To create a 20GB RAID 5 array for data files, we need four 10GB disks (three for data and one for parity). To create a 20GB RAID 5 array for the operating system, we need another four 10GB disks (three for data and one for parity). Therefore, the total number of disks needed is  $2 + 4 + 4 = 10$ . However, since we can use different RAID levels for different partitions on the same disk, we can optimize the disk usage by using only eight disks as follows: Disk 1: 10GB RAID 1 (log files) + 10GB RAID 5 (data files) Disk 2: 10GB RAID 1 (log files) + 10GB RAID 5 (data files) Disk 3: 10GB RAID 5 (data files) + 10GB RAID 5 (OS) Disk 4: 10GB RAID 5 (data files) + 10GB RAID 5 (OS) Disk 5: 10GB RAID 5 (parity for data files) + 10GB RAID 5 (OS) Disk 6: 10GB RAID 5 (OS) + unused space Disk 7: 10GB RAID 5 (parity for OS) + unused space Disk 8: unused space References: [https://en.wikipedia.org/wiki/Standard\\_RAID\\_levels](https://en.wikipedia.org/wiki/Standard_RAID_levels)

**NEW QUESTION 60**

A server administrator implemented a new backup solution and needs to configure backup methods for remote sites. These remote sites have low bandwidth and backups must not interfere with the network during normal business hours. Which of the following methods can be used to meet these requirements? (Select two).

- A. Open file
- B. Archive
- C. Cloud
- D. Snapshot
- E. Differential
- F. Synthetic full

**Answer:** BE

**Explanation:**

Archive is a method of storing historical data that is not frequently accessed or modified. Archive can reduce the amount of data that needs to be backed up and save bandwidth and storage space. Differential is a method of backing up only the data that has changed since the last full backup. Differential can also save bandwidth and storage space, as well as speed up the backup process.

References:

CompTIA Server+ Certification Exam Objectives1, page 12

Server Management: Server Hardware Installation and Management2, Module 2, Lesson 5

**NEW QUESTION 64**

A technician recently applied a critical OS patch to a working sever. 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 host firewall rule.
- D. Roll back the applied patch.

**Answer:** C

**Explanation:**

A host firewall is a software that controls the incoming and outgoing network traffic on a server based on predefined rules and filters. It can block or allow certain ports, protocols, or addresses that are used for communication with other servers or devices. If a server is unable to connect to another server after applying a patch, it is possible that the patch changed or added a firewall rule that prevents the connection. The administrator should check the host firewall rule and modify it if necessary to restore connectivity. Verified References: [Host firewall], [Network connection]

**NEW QUESTION 69**

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

**Answer:** A

**Explanation:**

This is the best definition of overprovisioning in a virtual server environment because it means allocating more CPU, memory, disk, or network resources to the virtual machines than what is actually available on the physical host. This can lead to performance issues and resource contention.

References: <https://www.hpe.com/us/en/insights/articles/10-virtualization-mistakes-everyone-makes-1808.html>

**NEW QUESTION 70**

Two developers are working together on a project, and they have built out a set of shared servers that both developers can access over the internet. Which of the following cloud models is this an example of?

- A. Hybrid
- B. Public
- C. Private
- D. Community

**Answer:** B

**Explanation:**

A public cloud is a cloud model that provides shared resources and services over the internet to multiple users or organizations. The cloud provider owns and manages the infrastructure and charges users based on their usage or subscription. A public cloud can offer scalability, flexibility, and cost-efficiency for users who need access to various applications and data without investing in their own hardware or software. Verified References: [Public cloud], [Cloud model]

**NEW QUESTION 74**

Which of the following symbols is used to write a text description per line within a PowerShell script?

- A. %
- B. @
- C. &
- D. #

**Answer:** D

**Explanation:**

The # symbol is used to write a text description per line within a PowerShell script. A text description is also known as a comment, which is a line of code that is ignored by the PowerShell interpreter and serves as documentation or explanation for human readers. The # symbol indicates that everything following it on the same line is a comment and not part of the script commands or expressions. For example:

This is a comment in PowerShell: Write-Host "Hello World" # This command prints Hello World to the console

References: CompTIA Server+ Certification Exam Objectives, Domain 6.0: Troubleshooting, Objective 6.3: Given a scenario, troubleshoot scripting errors using PowerShell commands.

**NEW QUESTION 76**

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.

**Answer:** C

**Explanation:**

Since both drives are in different RAID 1 arrays and both are hot-swappable, the most efficient step is to replace both failed drives at the same time. This can minimize the downtime and avoid unnecessary reboots. RAID 1 provides mirroring, which means that data is duplicated on both drives in the array. Therefore, replacing one drive will not affect the data on the other drive or the functionality of the array.

References: [https://en.wikipedia.org/wiki/Standard\\_RAID\\_levels#RAID\\_1](https://en.wikipedia.org/wiki/Standard_RAID_levels#RAID_1)

**NEW QUESTION 81**

A server administrator notices the /var/log/audit/audit.log file on a Linux server is rotating too frequently. The administrator would like to decrease the number of times the log rotates without losing any of the information in the logs. Which of the following should the administrator configure?

- A. increase the audit
- B. log file size in the appropriate configuration file.
- C. Decrease the duration of the log rotate cycle for the audit
- D. log file.
- E. Remove the log rotate directive from the audit .log file configuration.
- F. Move the audit
- G. log files to a remote syslog server.

**Answer:** A

**Explanation:**

The audit.log file is a file that records security-related events on a Linux server, such as user login, file access, and system commands. The logrotate utility is a tool that rotates, compresses, and deletes old log files based on certain criteria, such as size, time, or frequency. To decrease the number of times the log rotates without losing any information, the administrator should increase the audit.log file size in the appropriate configuration file, such as /etc/logrotate.conf or /etc/logrotate.d/auditd. Verified References: [audit.log], [logrotate]

**NEW QUESTION 83**

Which of the following should a technician verify FIRST before decommissioning and wiping a file server?

- A. The media destruction method
- B. The recycling process
- C. Asset management documentation
- D. Non-utilization

**Answer:** D

**Explanation:**

The first thing that a technician should verify before decommissioning and wiping a file server is non-utilization, which means that no one is using or accessing the server or its data. This can be done by checking logs, monitoring network traffic, or contacting users or stakeholders. Non-utilization ensures that decommissioning and wiping will not cause any data loss or disruption to business operations. Verified References: [Server Decommissioning Checklist]

**NEW QUESTION 86**

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. A VPN
- D. A virtual switch
- E. A vNIC

**Answer:** D

**Explanation:**

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 VLANs.

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.

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.

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 VLAN.

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.

**NEW QUESTION 91**

A user can successfully connect to a database server from a home office but is unable to access it from a hotel room. Which of the following authentication methods is most likely configured?

- A. Delegation
- B. Role-based
- C. Rule-based
- D. Scope-based

**Answer:** D

**Explanation:**

Scope-based authentication is a method of restricting access to resources based on the location, network, or device of the user. It can be used to prevent unauthorized access from outside the organization's network or from untrusted devices. In this case, the user can connect to the database server from the home office, which is likely within the scope of the authentication policy, but not from the hotel room, which is outside the scope.

References:

CompTIA Server+ Certification Exam Objectives<sup>1</sup>, page 15 CompTIA Server+: Authentication & Authorization<sup>2</sup>

**NEW QUESTION 94**

A remote, embedded IoT server is having a Linux OS upgrade installed. Which of the following is the best method to stage the new media for the default boot device of the server?

- A. Copy and send an SSD to the site.
- B. Copy and send a DVD to the site.
- C. Copy and send a SATA drive to the site.
- D. Copy and send a microSD card to the site.

**Answer:** D

**Explanation:**

A microSD card is the best method to stage the new media for the default boot device of a remote embedded IoT server that is having a Linux OS upgrade installed. A microSD card is a small and portable storage device that can store large amounts of data. It can be easily inserted into the slot of an embedded IoT server, which is a small and low-power device that performs specific tasks and connects to other devices over a network. A microSD card can also be formatted with different file systems, such as FAT32 or ext4, which are compatible with Linux OS. References: CompTIA Server+ Certification Exam Objectives, Domain 4.0: Networking, Objective 4.3: Given a scenario, configure servers for IoT applications.

**NEW QUESTION 99**

A company needs a media server set up that provides the highest availability with a minimum requirement of at least 10TB. The company purchased five HDDs, each with a 4TB capacity. Which of the options would provide the highest fault tolerance and meet the requirements?

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

**Answer:** C

**Explanation:**

RAID 6 is a RAID level that uses disk striping with two parity blocks distributed across all member disks. It can tolerate the failure of up to two disks without losing any data. RAID 6 can provide a minimum of 10TB of usable storage space with five 4TB disks, as the formula for calculating the RAID 6 capacity is  $(n-2) \times S_{min}$ , where  $n$  is the number of disks and  $S_{min}$  is the smallest disk size. In this case, the RAID 6 capacity is  $(5-2) \times 4TB = 12TB$ . References:

? CompTIA Server+ Certification Exam Objectives<sup>1</sup>, page 8

? RAID Levels and Types Explained: Advantages and Disadvantages<sup>2</sup>

? RAID Levels & Fault Tolerance<sup>3</sup>

**NEW QUESTION 101**

An administrator is helping to replicate a large amount of data between two Windows servers. The administrator is unsure how much data has already been transferred. Which of the following will BEST ensure all the data is copied consistently?

- A. rsync
- B. copy
- C. scp
- D. robocopy

**Answer:** D

**Explanation:**

Robocopy (Robust File Copy) is a command-line tool that can copy files and folders between Windows servers or computers. It has many features and options that can ensure all the data is copied consistently, such as retrying failed copies, resuming interrupted copies, copying permissions and attributes, mirroring source and destination directories, and logging the copy progress and results. Verified References: [Robocopy], [File copy]

**NEW QUESTION 104**

A technician is deploying a single server to monitor and record security cameras at a remote site, which of the following architecture types should be used to minimize cost?

- A. Virtual
- B. Blade
- C. Tower
- D. Rack mount

**Answer:** C

**Explanation:**

A tower server is a type of server architecture that is best suited to minimize cost when deploying a single server to monitor and record the security cameras at a remote site. A tower server is a standalone server that has a similar form factor and design as a desktop computer. It does not require any special mounting equipment or rack space and can be placed on or under a desk or table. A tower server is suitable for small businesses or remote offices that need only one or few servers for basic tasks such as file sharing, print serving, or security monitoring. A tower server is usually cheaper and easier to maintain than other types of servers, but it may have lower performance, scalability, and redundancy features. A virtual server is a type of server architecture that involves creating and running one or more virtual machines on a physical host using a hypervisor such as Hyper-V or VMware. A virtual server can reduce hardware costs and improve flexibility and efficiency, but it requires additional software licenses and management tools. A blade server is a type of server architecture that involves inserting multiple thin servers called blades into a chassis that provides power, cooling, network, and management features. A blade server can improve performance, density, and scalability, but it requires more initial investment and specialized equipment. A rack mount server is a type of server architecture that involves mounting one or more servers into standardized frames called racks that provide power, cooling, network, and security features.

**NEW QUESTION 106**

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

**Answer:** C

**Explanation:**

Encryption at rest is a type of encryption methodology that would most likely be used to ensure encrypted data cannot be retrieved if a device is stolen. Encryption at rest is a process of encrypting stored data on a device such as a hard drive, SSD, USB flash drive, or mobile device. This way, if the device is lost or stolen, the data cannot be accessed without the encryption key or password. Encryption at rest can be implemented using software tools such as BitLocker on Windows or FileVault on Mac OS, or hardware features such as self-encrypting drives or Trusted Platform Module chips. End-to-end encryption is a type of encryption methodology that ensures encrypted data cannot be intercepted or modified by third parties during transmission over a network. Encryption in transit is a type of encryption methodology that protects encrypted data while it is moving from one location to another over a network. Public key encryption is a type of encryption algorithm that uses a pair of keys: a public key that can be shared with anyone and a private key that is kept secret by the owner. References: <https://www.howtogeek.com/196541/bitlocker-101-what-it-is-how-it-works-and-how-to-use-it/> <https://www.howtogeek.com/443611/how-to-encrypt-your-macs-system-drive-removable-devices-and-individual-files/> <https://www.howtogeek.com/428483/what-is-end-to-end-encryption-and-why-does-it-matter/> <https://www.howtogeek.com/195877/what-is-encryption-and-how-does-it-work/>

**NEW QUESTION 111**

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

**Answer: B**

**Explanation:**

The administrator must encrypt the data at rest to ensure data on the SAN is not compromised if it is leaked. Data at rest refers to data that is stored on a device or a medium, such as a hard drive, a flash drive, or a SAN (Storage Area Network). Data at rest can be leaked if the device or the medium is lost, stolen, or accessed by unauthorized parties. Encrypting data at rest means applying an algorithm that transforms the data into an unreadable format that can only be decrypted with a key. Encryption protects data at rest from being exposed or misused by attackers who may obtain the device or the medium.

**NEW QUESTION 116**

After installing a new file server, a technician notices the read times for accessing the same file are slower than the read times for other file servers. Which of the following is the first step the technician should take?

- A. Add more memory.
- B. Check if the cache is turned on.
- C. Install faster hard drives.
- D. Enable link aggregation.

**Answer: B**

**Explanation:**

The cache is a temporary storage area that holds frequently accessed data or instructions for faster retrieval. The cache can improve the read times for accessing files by reducing the need to access the hard drive, which is slower than the cache memory<sup>1</sup>. Therefore, the first step the technician should take is to check if the cache is turned on for the new file server. If the cache is turned off, the technician should enable it and see if the read times improve. The other options are incorrect because they are not the first steps to take. Adding more memory, installing faster hard drives, or enabling link aggregation are possible ways to improve the performance of the file server, but they are more costly and time-consuming than checking the cache. Moreover, they may not address the root cause of the problem if the cache is turned off.

**NEW QUESTION 118**

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

**Answer: C**

**Explanation:**

This is the most likely misconfigured setting because DNS is the service that resolves hostnames to IP addresses and vice versa. If the DNS server is incorrect or unreachable, the administrator will not be able to log on to the server with a domain account because the server will not be able to authenticate with the domain controller.

References: <https://docs.microsoft.com/en-us/troubleshoot/windows-server/networking/dns-troubleshooting>

**NEW QUESTION 122**

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

**Answer: A**

**Explanation:**

Reference: <https://www.druva.com/blog/understanding-rpo-and-rto/>

The Recovery Time Objective (RTO) is the maximum amount of time that a company agrees to tolerate in the event of a disaster before restoring its normal operations. The RTO is based on the business impact analysis (BIA) and the criticality of the processes and data involved. The RTO helps determine the backup and recovery strategies and resources needed to minimize downtime and data loss.

Reference: <https://www.ibm.com/cloud/learn/recovery-time-objective>

**NEW QUESTION 127**

A user has been unable to authenticate to the company's external, web-based database after clicking a link in an email that required the user to change the account password. Which of the following steps should the company take next?

- A. Disable the user's account and inform the security team.
- B. Create a new log-in to the external database.
- C. Ask the user to use the link again to reset the password.
- D. Reset the user's password and ask the user to log in again.

**Answer:** A

**Explanation:**

The user has likely fallen victim to a phishing scam, which is a fraudulent attempt to obtain sensitive information, such as passwords, by disguising as a legitimate entity. The link in the email that required the user to change the account password was probably a fake website that mimicked the company's external database, and captured the user's credentials when they entered them. This could compromise the security and integrity of the company's data, as well as the user's identity and privacy<sup>12</sup>.

The company should take immediate action to prevent further damage and investigate the incident. The first step is to disable the user's account and inform the security team. Disabling the user's account can prevent unauthorized access to the external database by the attackers, who may use the stolen credentials to log in and manipulate or steal data. Informing the security team can alert them of the breach and allow them to take appropriate measures, such as scanning for malware, changing passwords, notifying other users, and reporting the incident<sup>34</sup>.

**NEW QUESTION 131**

A server administrator deployed a new product that uses a non-standard port for web access on port 8443. However, users are unable to access the new application. The server administrator checks firewall rules and determines 8443 is allowed. Which of the following is most likely the cause of the issue?

- A. Intrusion detection is blocking the port.
- B. The new application's DNS entry is incorrect.
- C. The application should be changed to use port 443.
- D. The core switch has a network issue.

**Answer:** B

**Explanation:**

A DNS entry is a record that maps a domain name to an IP address. If the DNS entry for the new application is incorrect, users will not be able to resolve the domain name to the correct IP address and port number. This will prevent them from accessing the application, even if the firewall rules allow port 8443. To fix this issue, the server administrator should verify and update the DNS entry for the new application.

References: CompTIA Server+ Study Guide, Chapter 6: Networking, page 230.

**NEW QUESTION 132**

A staff member who is monitoring a data center reports one rack is experiencing higher temperatures than the racks next to it, despite the hardware in each rack being the same. Which of the following actions would MOST likely remediate the heat issue?

- A. Installing blanking panels in all the empty rack spaces
- B. installing an additional POU and spreading out the power cables
- C. Installing servers on the shelves instead of sliding rails
- D. installing front bezels on all the server's in the rack

**Answer:** A

**Explanation:**

Blanking panels are metal or plastic plates that are installed in the empty spaces of a rack to prevent hot air from recirculating back to the front of the rack. This can improve the airflow and cooling efficiency of the rack and reduce the heat generated by the servers. Verified References: [Blanking panel], [Rack cooling]

**NEW QUESTION 137**

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

**Answer:** B

**Explanation:**

A disaster recovery plan would be useful when trying to restore IT infrastructure operations after a non-planned interruption. A disaster recovery plan is a document that outlines the steps and procedures to recover from a major disruption of IT services caused by natural or man-made disasters, such as fire, flood, earthquake, cyberattack, etc. A disaster recovery plan typically includes:

- ? A list of critical IT assets and resources that need to be protected and restored
- ? A list of roles and responsibilities of IT staff and stakeholders involved in the recovery process
- ? A list of backup and recovery strategies and tools for data, applications, servers, networks, etc.
- ? A list of communication channels and methods for notifying users, customers, vendors, etc.
- ? A list of testing and validation methods for ensuring the functionality and integrity of restored systems
- ? A list of metrics and criteria for measuring the effectiveness and efficiency of the recovery process

A disaster recovery plan helps IT organizations to minimize downtime, data loss, and financial impact of a disaster, as well as to resume normal operations as quickly as possible.

**NEW QUESTION 138**

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

**Answer:** B

**Explanation:**

The tool that the administrator should use to check for unnecessary running services across 12 servers is a port scanner. A port scanner is a tool that scans a network device for open ports and identifies the services or applications that are running on those ports. A port scanner can help detect any unauthorized or unwanted services that may pose a security risk or consume network resources. A port scanner can also help troubleshoot network connectivity issues or verify firewall rules.

Reference: <https://www.getsafeonline.org/business/articles/unnecessary-services/>

**NEW QUESTION 143**

A technician is troubleshooting a server issue. The technician has determined several possible causes of the issue and has identified various solutions. Which of the following should the technician do next?

- A. Consult internet forums to determine which is the most common cause and deploy only that solution.
- B. Test each solution individually to determine the root cause, rolling back the changes in between each test.
- C. Implement the shortest solution first to identify the issue and minimize downtime.
- D. Test each solution in succession and restore the server from the latest snapshot.

**Answer:** B

**Explanation:**

According to the CompTIA troubleshooting methodology, the fourth step is to establish a plan of action to resolve the problem and implement the solution<sup>1</sup>. The best practice is to test each solution individually to determine the root cause, rolling back the changes in between each test. This way, the technician can isolate the cause and avoid introducing new problems or making the situation worse. Testing each solution in succession and restoring the server from the latest snapshot (D) is not a good option because it may not identify the root cause and may overwrite important data. Implementing the shortest solution first to identify the issue and minimize downtime © is also not a good option because it may not solve the problem or may create new issues. Consulting internet forums to determine which is the most common cause and deploy only that solution (A) is not a good option because it may not apply to the specific situation or may be outdated or inaccurate

**NEW QUESTION 146**

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

**Answer:** D

**Explanation:**

The server type that would benefit most from the use of a load balancer is web server. A web server is a server that hosts web applications or websites and responds to requests from web browsers or clients. A load balancer is a device or software that distributes network traffic across multiple servers based on various criteria, such as availability, capacity, or performance. A load balancer can improve the scalability, reliability, and performance of web servers by balancing the workload and preventing any single server from being overloaded or unavailable.

Reference:

<https://www.dnsstuff.com/what-is-server-load-balancing>

**NEW QUESTION 150**

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.

**Answer:** C

**Explanation:**

The first action that the administrator should take before swapping out the GPU card inside a server is to shut down the server. This is to ensure that the server is not running any processes that might be using the GPU card, and to prevent any damage to the hardware or data loss due to sudden power loss. Shutting down the server also reduces the risk of electrostatic discharge (ESD) that might harm the components. Reference: <https://pcgearhead.com/installing-a-new-gpu/>

**NEW QUESTION 152**

A technician is creating a network share that will be used across both Unix and Windows clients at the same time. Users need read and write access to the files. Which of the following would be BEST for the technician to deploy?

- A. iSCSI
- B. CIFS
- C. HTTPS
- D. DAS

**Answer:**

B

**Explanation:**

CIFS (Common Internet File System) is a protocol that allows file sharing across different operating systems, such as Unix and Windows. It supports read and write access to files and folders on a network share. It is also known as SMB (Server Message Block). Verified References: [CIFS], [File sharing]

**NEW QUESTION 156**

A human resources analyst is attempting to email the records for new employees to an outside payroll company. Each time the analyst sends an email containing employee records, the email is rejected with an error message. Other emails outside the company are sent correctly. Which of the following is MOST likely generating the error?

- A. DHCP configuration
- B. Firewall rules
- C. DLP software
- D. Intrusion detection system

**Answer:** C

**Explanation:**

DLP (Data Loss Prevention) software is a type of security software that monitors and controls the transfer of sensitive or confidential data outside the organization. DLP software can prevent data breaches, data leaks, or data theft by blocking, encrypting, or alerting on unauthorized data transfers. DLP software can be applied to various channels, such as email, web, cloud, or removable devices.

In this scenario, the human resources analyst is attempting to email the records for new employees to an outside payroll company. The records for new employees may contain sensitive or confidential data, such as personal information, tax information, or bank account information. The DLP software may detect this data and block the email from being sent outside the company, as it may violate the company's data protection policy or regulations. The DLP software may also generate an error message to inform the analyst of the reason for the rejection.

**NEW QUESTION 160**

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

**Answer:** B

**Explanation:**

The technician should use port 22 to configure a server that requires secure remote access. Port 22 is the default port for Secure Shell (SSH), which is a protocol that

allows secure remote login and command execution over a network connection using a command-line interface (CLI). SSH encrypts both the authentication and data transmission between the client and the server, preventing eavesdropping, tampering, or spoofing. SSH can be used to perform various tasks on a server remotely, such as configuration, administration, maintenance, troubleshooting, etc.

**NEW QUESTION 162**

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.

**Answer:** C

**Explanation:**

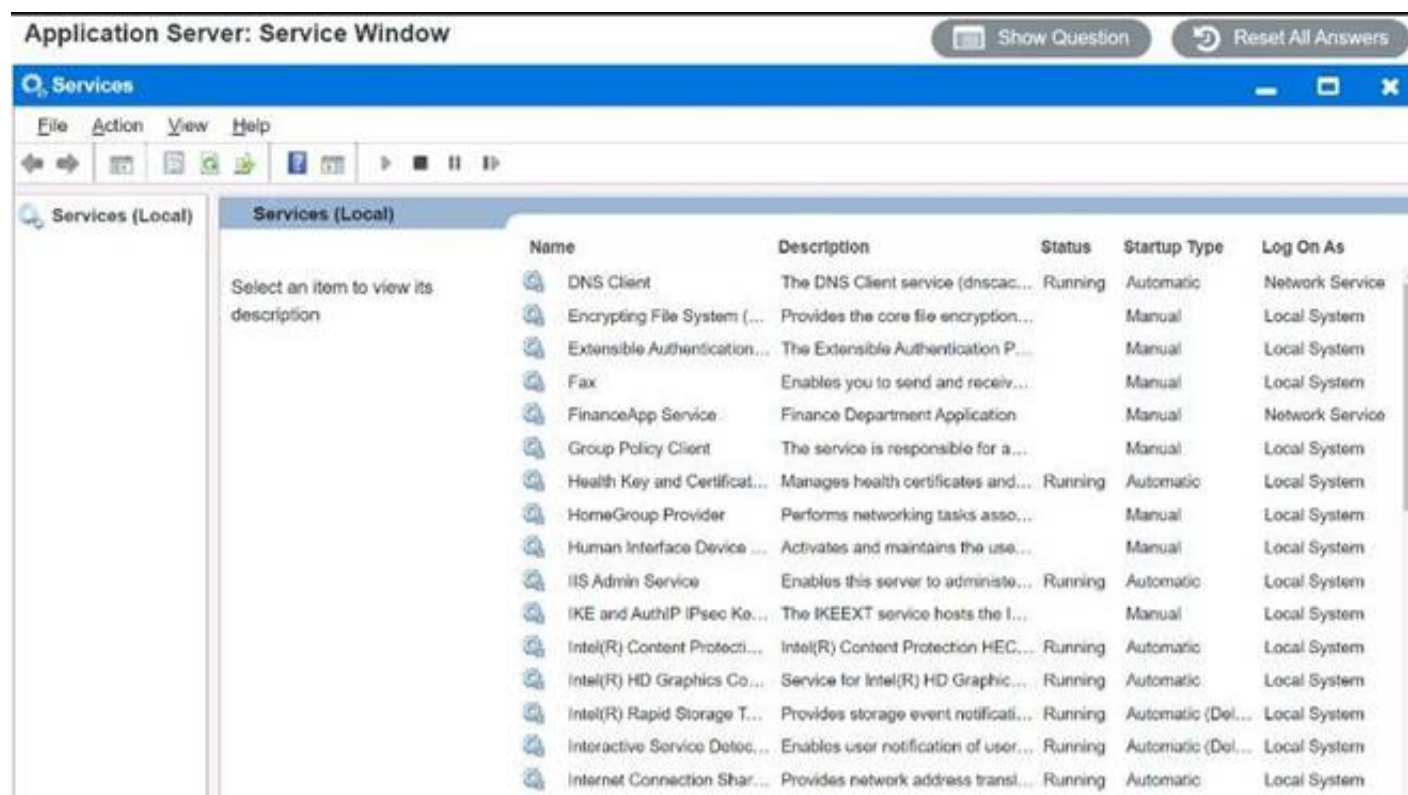
This is the best way to proceed with troubleshooting when the technician has several possible solutions to a reported server issue. Making one change at a time and testing allows the technician to isolate the cause and effect of each solution and determine which one works best. It also helps to avoid introducing new problems or complicating existing ones by making multiple changes at once. Determining whether there is a common element in the symptoms causing multiple problems is a good step to perform before identifying possible solutions, but not after. Performing a root cause analysis is a good step to perform after resolving the issue, but not during. Documenting the findings, actions, and outcomes throughout the process is a good practice to follow at every step of troubleshooting, but not a specific way to proceed with testing possible solutions. References: <https://www.howtogeek.com/174288/how-to-tell-if-your-computer-is-overheating-and-what-to-do-about-it/><https://www.howtogeek.com/428483/what-is-end-to-end-encryption-and-why-does-it-matter/>

**NEW QUESTION 167****SIMULATION**

Users report that the FinanceApp software is not running, and they need immediate access. Issues with the FinanceApp software occur every week after the IT team completes server system updates. The users, however, do not want to contact the help desk every time the issue occurs. The users also report the new MarketApp software is not usable when it crashes, which can cause significant downtime. The technician who restarted the MarketApp software noticed it is running under a test account, which is a likely cause of the crashes.

**INSTRUCTIONS**

Using the Services menu provided, modify the appropriate application services to remedy the stated issues.



- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

FinanceApp software is running as a service named “FinanceApp Service”. The service description says “Provides financial data and calculations for the FinanceApp software”. The service status is “Stopped”, which means that the service is not running and the software is not functional. The service startup type is “Manual”, which means that the service needs to be started manually by the user or the administrator. The service log on as is “Local System”, which means that the service runs under a predefined local account that has extensive privileges on the local computer.

To fix the issue with the FinanceApp software, you need to do two things:

? First, you need to start the service, so that the software can run. To do this, you can right-click on the service name and select “Start” from the menu.

Alternatively, you can select the service name and click on the “Start” button on the toolbar. You should see a message saying that the service has started successfully.

? Second, you need to change the service startup type, so that the service can start automatically every time the server boots up. This way, you don’t have to contact the help desk every time the issue occurs. To do this, you can right-click on the service name and select “Properties” from the menu. Alternatively, you can select the service name and click on the “Properties” button on the toolbar. You should see a window with several tabs and options. On the “General” tab, under “Startup type”, you can select “Automatic” from the drop-down list. Then, click on “OK” to save your changes.

By doing these two steps, you should be able to use the FinanceApp software without any problems.

The MarketApp software is running as a service named “MarketApp Service”. The service description says “Provides market data and analysis for the MarketApp software”. The service status is “Running”, which means that the service is running and the software is functional. However, as you reported, the software may crash sometimes, which can cause significant downtime. The service startup type is “Automatic”, which means that the service starts automatically every time the server boots up. The service log on as is “TestAccount”, which is a test account that was probably used for development or testing purposes.

To fix the issue with the MarketApp software, you need to do one thing:

? You need to change the service log on as, so that the service runs under a proper account that has sufficient permissions and security settings for production use. To do this, you can right-click on the service name and select “Properties” from the menu. Alternatively, you can select the service name and click on the “Properties” button on the toolbar. You should see a window with several tabs and options. On the “Log On” tab, under “Log on as”, you can select either “Local System account” or “This account”. If you choose “Local System account”, then the service will run under a predefined local account that has extensive privileges on the local computer. If you choose “This account”, then you will need to enter a valid username and password for an account that has appropriate permissions and security settings for running the service. You may need to consult with your IT team or your software vendor to determine which option is best for your situation. Then, click on “OK” to save your changes.

**NEW QUESTION 168**

An administrator has deployed a new virtual server from a template. After confirming access to the subnet’s gateway, the administrator is unable to log on with the domain credentials. Which of the following is the most likely cause of the issue?

- A. The server has not been joined to the domain.
- B. An IP address has not been assigned to the server.
- C. The server requires a reboot to complete the deployment process.
- D. The domain credentials are invalid.

**Answer: A**

**Explanation:**

The most likely cause of the issue is that the server has not been joined to the domain. A domain is a logical group of computers and devices that share a common directory service and security policy. A domain controller is a server that manages the domain and authenticates users and computers that want to access domain resources. To log on with domain credentials, a server must be joined to the domain and registered in the directory service. If a server has not been joined to the domain, it will not be recognized or authorized by the domain controller.

References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 4, Lesson 4.3, Objective 4.3

**NEW QUESTION 172**

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

**Answer:** D

**Explanation:**

The live failover testing method is the most likely one that the server administrator is performing when creating a downtime scenario and taking the necessary steps. A live failover test involves switching from the primary system to the secondary system (or backup site) in a real environment, without any simulation or preparation. A live failover test can evaluate the effectiveness and readiness of the disaster recovery plan, but it also carries a high risk of data loss, corruption, or disruption. Reference: <https://www.ibm.com/cloud/learn/disaster-recovery-testing>

**NEW QUESTION 174**

A technician is working on a Linux server and is trying to access another server over the network. The technician gets a server not found message when trying to execute `ping servername` but no error messages when using `pingservername.domain.com`. Which of the following should the technician do to resolve the error?

- A. Configure the domain search variable
- B. Change the permissions on `/etc/resolv.conf`
- C. `/etc/resolv.conf`
- D. Configure the DNS address
- E. Modify `/etc/nsswitch.conf`
- F. `/etc/nsswitch.conf`

**Answer:** A

**Explanation:**

The domain search variable is used to specify a list of domains that are appended to a hostname when resolving it. If the servername is not fully qualified, the resolver will try each domain in the list until it finds a match or fails. By configuring the domain search variable, the technician can avoid typing the full domain name every time they want to ping a server. Verified References: [How to configure DNS suffixes on Linux systems]

**NEW QUESTION 177**

Which of the following, if properly configured, would prevent a user from installing an OS on a server? (Select TWO).

- A. Administrator password
- B. Group Policy Object
- C. Root password
- D. SELinux
- E. Bootloader password
- F. BIOS/UEFI password

**Answer:** EF

**Explanation:**

These are two methods that can prevent a user from installing an OS on a server if properly configured. A bootloader password is a password that protects the bootloader from unauthorized access or modification. The bootloader is a program that loads the operating system into memory when the system boots up. If a user does not know the bootloader password, they cannot change the boot order or boot from another device such as a CD-ROM or USB drive that contains an OS installation media. A BIOS/UEFI password is a password that protects the BIOS (Basic Input Output System) or UEFI (Unified Extensible Firmware Interface) from unauthorized access or modification. The BIOS or UEFI is a firmware that initializes and configures the hardware components of the system before loading

**NEW QUESTION 178**

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

**Answer:** D

**Explanation:**

The DR testing scenario that is described as verbally walking through each step of the DR plan in the context of a meeting is tabletop. A tabletop test is a type of disaster recovery (DR) test that involves discussing and reviewing the DR plan with key stakeholders and participants in a simulated scenario. A tabletop test does not involve any actual execution of the DR plan or any disruption of the normal operations. A tabletop test can help identify gaps, issues, or inconsistencies in the DR plan and improve communication and coordination among the DR team members.

**NEW QUESTION 181**

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

**Answer:** B

**Explanation:**

The cluster feature that will allow Ann to complete her task is failback. A cluster is a group of servers that work together to provide high availability, scalability, and

load balancing for applications or services. A cluster can have different nodes or members that have different roles or states. An active node is a node that is currently running an application or service and serving requests from clients. A passive node is a node that is on standby and ready to take over if the active node fails. A failover is a process of switching from a failed or unavailable node to another node in a cluster. A failback is a process of switching back from a failover node to the original node after it becomes available again. Failback can be automatic or manual depending on the cluster configuration.

**NEW QUESTION 186**

An administrator is troubleshooting a server that is rebooting and crashing. The administrator notices that the server is making sounds that are louder than usual. Upon closer inspection, the administrator discovers that the noises are coming from the front of the chassis. Which of the following is the most likely reason for this behavior?

- A. One of the fans has failed.
- B. The power supply has failed.
- C. The RAM is malfunctioning.
- D. The CPU is overheating.

**Answer:** A

**Explanation:**

A server has multiple fans inside the chassis to cool down the components and prevent overheating. If one of the fans fails, it can cause the server to reboot and crash due to thermal issues. A failed fan can also make loud noises due to friction or vibration. The administrator should check the fans and clean them from dust and debris, or replace them if they are damaged<sup>12</sup>.

References = 1: It's Too Loud! 3 Solutions to Remedy Server Noise - Computerware Blog | DC Metro | Computerware Blog(<https://www.cwit.com/blog/it-s-too-loud-3-solutions-to-remedy-server-noise>) 2: What factors affect the noise level of a server? - Server Fault(<https://serverfault.com/questions/430550/what-factors-affect-the-noise-level-of-a-server>)

**NEW QUESTION 191**

An administrator is troubleshooting an application performance issue on a virtual server with two vCPUs. The application performance logs indicate CPU contention. The administrator adds more vCPU cores to the VM, yet the issue persists. Which of the following is the most likely reason for this issue?

- A. The server has high page utilization.
- B. The server has high disk latency.
- C. The application is single-threaded.
- D. The application cannot be virtualized.

**Answer:** C

**Explanation:**

A single-threaded application is an application that can only execute one task or process at a time. A single-threaded application can only utilize one CPU core, regardless of how many cores are available or assigned to the virtual machine. Therefore, adding more vCPU cores to the VM will not improve the performance of the application, as it will still be limited by the speed and capacity of one core<sup>12</sup>.

To troubleshoot this issue, the administrator should check if the application is single-threaded or multi-threaded. This can be done by using tools such as Task Manager, Performance Monitor, or Process Explorer on Windows, or top, htop, or ps on Linux<sup>34</sup>. If the application is single-threaded, the administrator should consider the following options:

- ? Reduce the number of vCPU cores on the VM to match the number of threads that the application can use. This can help avoid CPU contention and co-stop issues that may arise from having too many vCPUs relative to the number of physical cores on the host<sup>5</sup>.
- ? Upgrade the physical CPU on the host to a faster or newer model that can provide higher clock speed and performance for the single core that the application uses.
- ? Optimize the application code or configuration to make it more efficient or multi-threaded, if possible. This can help the application take advantage of multiple cores and improve its performance.

**NEW QUESTION 193**

Which of the following licensing models is MOST appropriate for a data center that has a variable daily equipment count?

- A. Per site
- B. Per server
- C. Per user
- D. Per core

**Answer:** D

**Explanation:**

A per core licensing model is based on the number of processor cores in a server. This model is suitable for a data center that has a variable daily equipment count, as it allows for scaling up or down the number of cores as needed. A per core licensing model also provides better performance and efficiency than other models. Verified References: [Per Core Licensing and Basic Definitions]

**NEW QUESTION 198**

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

**Answer:** D

**Explanation:**

netstat is a command-line tool that displays network connections, routing tables, interface statistics, and more. The -a option shows all listening and non-listening sockets on the server. This can help check the open ports on a server and identify any unwanted or malicious connections. References: <https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/netstat>

**NEW QUESTION 201**

An administrator is working locally in a data center with multiple server racks. Which of the following is the best low-cost option to connect to any server while on site?

- A. Crash cart
- B. IPKVM
- C. Remote console access
- D. IPMI

**Answer:** A

**Explanation:**

A crash cart is the best low-cost option to connect to any server while on site in a data center with multiple server racks. A crash cart is a mobile unit that contains a monitor, a keyboard, a mouse, and cables that can be plugged into any server for direct access and control. A crash cart can be used for troubleshooting, maintenance, or configuration of servers without requiring remote access or network connectivity. A crash cart is also easy to move around and store in a data center. References: [CompTIA Server+ Certification Exam Objectives], Domain 2.0: Hardware, Objective 2.4: Given a scenario involving server management issues (e.g., remote access), troubleshoot using appropriate tools.

**NEW QUESTION 203**

Which of the following would a systems administrator most likely implement to encrypt data in transit for remote administration?

- A. Telnet
- B. SSH
- C. TFTP
- D. rlogin

**Answer:** B

**Explanation:**

SSH (Secure Shell) is a protocol that would most likely be implemented to encrypt data in transit for remote administration. SSH provides secure communication between two devices over an unsecured network by using public-key cryptography and symmetric encryption. SSH can be used to remotely execute commands, transfer files, or tunnel other protocols. Telnet, TFTP, and rlogin are protocols that do not encrypt data in transit and are considered insecure for remote administration. References: [CompTIA Server+ Certification Exam Objectives], Domain 2.0: Networking, Objective 2.4: Given a scenario involving network security/access methods, implement an appropriate solution.

**NEW QUESTION 206**

A software developer is unable to reach an internal website. The developer's attempt to ping the FQDN returns the following IP address: 104.18.17.32. Which of the following is the most likely reason for this result?

- A. The NIC is set to DHCP.
- B. The default gateway is misconfigured.
- C. The primary DNS server is 8.8.8.8.
- D. There is a manual entry in the hosts file.

**Answer:** D

**Explanation:**

The most likely reason for this result is that there is a manual entry in the hosts file that maps the FQDN to an incorrect IP address (104.18.17.32). The hosts file is a text file that contains mappings of hostnames or domain names to IP addresses, which are used by the operating system to resolve names before querying DNS servers on the network or internet. The hosts file can be used to override DNS settings or block access to certain websites by redirecting them to different IP addresses, such as localhost (127.0.0.1) or invalid addresses (0.0.0.0). If there is a manual entry in the hosts file that conflicts with DNS records, it can cause name resolution errors or connectivity issues. References: [CompTIA Server+ Certification Exam Objectives], Domain 4.0: Networking, Objective 4.1: Given a scenario, configure network settings for servers.

**NEW QUESTION 207**

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

**Answer:** A

**Explanation:**

The best course of action for the company is to upgrade the application package to fix the known vulnerability. A vulnerability is a weakness or flaw in an application that can be exploited by an attacker to compromise the security or functionality of the system. Upgrading the application package means installing a newer version of the application that has patched or resolved the vulnerability. This way, the company can prevent potential attacks that may exploit the vulnerability and cause damage or loss.

**NEW QUESTION 208**

A systems administrator recently installed a new virtual server. After completing the installation, the administrator was only able to reach a few of the servers on the network. While testing, the administrator discovered only servers that had similar IP addresses were reachable. Which of the following is the most likely cause

of the issue?

- A. The jumbo frames are not enabled.
- B. The subnet mask is incorrect.
- C. There is an IP address conflict.
- D. There is an improper DNS configuration.

**Answer:** B

**Explanation:**

A subnet mask is a number that distinguishes the network address and the host address within an IP address<sup>1</sup>. A subnet mask allows network traffic to understand IP addresses by splitting them into the network and host addresses. If the subnet mask is incorrect, the network traffic may not be able to determine the correct destination for the packets, and only reach some of the servers that have similar IP addresses. For example, if the new virtual server has an IP address of 192.168.1.100 and a subnet mask of 255.255.0.0, it can only communicate with servers that have IP addresses in the range of 192.168.0.0 to 192.168.255.255. To fix this issue, the systems administrator needs to check and correct the subnet mask of the new virtual server according to the network configuration.

**NEW QUESTION 211**

A server administrator is taking advantage of all the available bandwidth of the four NICs on the server. Which of the following NIC-teaming technologies should the server administrator utilize?

- A. Fail over
- B. Fault tolerance
- C. Load balancing
- D. Link aggregation

**Answer:** D

**Explanation:**

Link aggregation is a technique that combines multiple physical network links into one logical link with higher bandwidth and redundancy. It can take advantage of all the available bandwidth of the NICs (Network Interface Cards) on the server and provide load balancing and failover capabilities for network traffic. Verified References: [Link aggregation], [NIC]

**NEW QUESTION 213**

A system 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

**Answer:** A

**Explanation:**

The best way to limit an attacker from exploiting a zero-day vulnerability that is impacting a service enabled on a server OS is to install the latest patches. Patches are updates that fix bugs, improve security, or add features to software. Installing patches can help prevent attackers from exploiting known vulnerabilities that have been fixed by the software vendor. A zero-day vulnerability is a vulnerability that is unknown to the vendor or the public until it is exploited by an attacker. Therefore, installing patches as soon as they are available can reduce the window of opportunity for attackers to exploit zero-day vulnerabilities. Reference: <https://www.ibm.com/cloud/learn/patch-management>

**NEW QUESTION 214**

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 boot. 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.

**Answer:** A

**Explanation:**

If a server has two arrays on the failed server: a two-drive RAID 0 set for the OS, and an eight-drive RAID 10 set for data, then the most likely failure scenario that caused the server to crash and not boot is that a drive failed in the OS array. RAID 0 is a RAID configuration that stripes data across two or more drives without parity or redundancy. RAID 0 offers high performance but no fault tolerance. If one drive fails in RAID 0, all data is lost and the system cannot boot. RAID 10 is a RAID configuration that combines disk mirroring and disk striping with parity. RAID 10 offers high performance and fault tolerance. RAID 10 can tolerate up to one drive failure per mirrored pair without losing data or functionality. References: <https://www.technewstoday.com/what-is-a-raid-0/>  
<https://www.technewstoday.com/what-is-a-raid-10/>

**NEW QUESTION 216**

A technician is working on a Linux server. The customer has reported that files in the home directory are missing. The /etc/fstab file has the following entry:

```
nfsserver:/home /home nfs defaults 0 0
```

However, a `df -h /home` command returns the following information:

```
/dev/sda2 10G 1G 9G 10% /home
```

Which of the following should the technician attempt FIRST to resolve the issue?

- A. `mkdir /home`
- B. `umount nfsserver:/home`

C. rmdir nfsserver:/home/dev/sda2  
D. mount /home

**Answer: B**

**Explanation:**

The /etc/fstab file contains the information about the file systems that are mounted automatically at boot time or on demand. The entry nfsserver:/home /home nfs defaults 0 0 indicates that the /home directory on the local server is mounted from the /home directory on a remote server called nfsserver using the NFS protocol. However, the df -h /home command shows that the /home directory is actually mounted from a local partition /dev/sda2, which may not contain the user's files. This means that the NFS mount failed or was overridden by another mount. To resolve the issue, the technician should attempt to unmount the local partition using umount nfsserver:/home, which will detach the /home directory from /dev/sda2. Then, the technician should try to mount the NFS share again using mount /home, which will attach the /home directory to nfsserver:/home according to the /etc/fstab entry<sup>12</sup>. Creating a new directory (A) or removing an existing one © would not help, as they would not affect the mount point. Mounting /home (D) without unmounting it first would not work, as it would result in an error that the mount point is busy<sup>3</sup>. References: 1 <https://askubuntu.com/questions/374870/home-directory-not-being-created> 2 <https://www.techrepublic.com/article/how-to-properly-automount-a-drive-in-ubuntu-linux/> 3 <https://serverfault.com/questions/587855/cannot-find-home-directory-on-linux-server>

**NEW QUESTION 220**

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.

**Answer: A**

**Explanation:**

Monitors Network Traffic Reference:<https://www.fortinet.com/resources/cyberglossary/benefits-of-firewall>

**NEW QUESTION 225**

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\$

**Answer: B**

**Explanation:**

The administrator should use \\server\%username% to correctly map a script to a home directory for a user based on username. %username% is an environment variable that represents the current user's name on a Windows system. By using this variable in the path of the script, the administrator can dynamically map the script to the user's home directory on the server. For example, if the user's name is John, the script will be mapped to \\server\John.

Reference:

<https://social.technet.microsoft.com/Forums/windows/en-US/07cfc73-796d-48aa-96a9-08280a1ef25a/mapping-home-directory-with-username-variable?forum=w7itprogeneral>

**NEW QUESTION 229**

An administrator is troubleshooting a failed NIC in an application server. The server uses DHCP to get all IP configurations, and the server must use a specific IP address. The administrator replaces the NIC, but then the server begins to receive a different and incorrect IP address. Which of the following will enable the server to get the proper IP address?

- A. Modifying the MAC used on the DHCP reservation
- B. Updating the local hosts file with the correct IP address
- C. Modifying the WWNN used on the DHCP reservation
- D. Updating the NIC to use the correct WWNN

**Answer: A**

**Explanation:**

A DHCP reservation is a way to assign a specific IP address to a device based on its MAC address, which is a unique identifier for each network interface card (NIC). When the administrator replaced the NIC, the MAC address of the server changed, and the DHCP server no longer recognized it as the same device. Therefore, the DHCP server assigned a different IP address to the server, which was incorrect for the application. To fix this problem, the administrator needs to modify the DHCP reservation to use the new MAC address of the NIC, so that the server can get the proper IP address.

A WWNN (World Wide Node Name) is a unique identifier for a Fibre Channel node, which is a device that can communicate over a Fibre Channel network. A WWNN is not related to DHCP or IP addresses, and it is not used for DHCP reservations. Therefore, options B and D are incorrect.

Updating the local hosts file with the correct IP address (option C) is also incorrect, because it does not solve the problem of getting the correct IP address from the DHCP server. The hosts file is a local file that maps hostnames to IP addresses, and it is used to override DNS queries. However, it does not affect how the DHCP server assigns IP addresses to devices. Moreover, updating the hosts file manually on every device that needs to communicate with the server is not a scalable or efficient solution.

References:

- ? How to reserve IP Address in DHCP Server - Ask Ubuntu
- ? Static IP vs DHCP Reservation - The Tech Journal
- ? How to Configure DHCP Server Reservation in Windows ... - ITIngredients

**NEW QUESTION 231**

.....

## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### SK0-005 Practice Exam Features:

- \* SK0-005 Questions and Answers Updated Frequently
- \* SK0-005 Practice Questions Verified by Expert Senior Certified Staff
- \* SK0-005 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* SK0-005 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The SK0-005 Practice Test Here](#)**