

CompTIA

Exam Questions SK0-005

CompTIA Server+ Certification Exam



NEW QUESTION 1

An administrator needs to increase the size of an existing RAID 6 array that is running out of available space. Which of the following is the best way the administrator can perform this task?

- A. Replace all the array drives at once and then expand the array.
- B. Expand the array by changing the RAID level to 6.
- C. Expand the array by changing the RAID level to 10.
- D. Replace the array drives one at a time and then expand the array.

Answer: D

Explanation:

RAID 6 is a type of RAID that uses block-level striping with two parity blocks distributed across all member disks. It allows for two disk failures within the RAID set before any data is lost¹. A minimum of four disks is required to create RAID 6¹. To increase the size of an existing RAID 6 array, the administrator can replace the array drives one at a time with larger drives and then expand the array. This way, the data and parity are rebuilt on each new drive and the array remains operational during the process².

NEW QUESTION 2

A systems administrator needs to create a data volume out of four disks with the MOST redundancy. Which of the following is the BEST solution?

- A. RAID 0
- B. RAID 1
- C. RAID 5
- D. RAID 6

Answer: D

Explanation:

RAID 6 is a type of RAID level that uses two parity blocks to provide fault tolerance and redundancy for data storage. RAID 6 can withstand the failure of up to two disks in the array without losing any data. RAID 6 requires a minimum of four disks to operate, and it distributes the data and parity blocks across all the disks in the array. RAID 6 has a high write penalty, which means that it takes more time and resources to write data to the disks than to read data from them. However, RAID 6 offers a high level of data protection and reliability, which makes it suitable for applications that require high availability and durability¹.

RAID 1 provides redundancy and fault tolerance by mirroring the data from one disk to another disk. RAID 1 offers high read performance and data security, but it has low capacity and write performance. RAID 1 requires a minimum of two disks to operate, and it can only tolerate the failure of one disk in the array. If more than one disk fails, all the data in the array is lost².

RAID 5 provides redundancy and fault tolerance by using one parity block to store information that can be used to reconstruct the data in case of a disk failure. RAID 5 requires a minimum of three disks to operate, and it distributes the data and parity blocks across all the disks in the array. RAID 5 offers a balance between performance, capacity, and data protection, but it can only tolerate the failure of one disk in the array. If more than one disk fails, all the data in the array is lost². Therefore, among these options, RAID 6 is the best solution for creating a data volume out of four disks with the most redundancy.

NEW QUESTION 3

A data center employee shows a driver's license to enter the facility. Once the employee enters, the door immediately closes and locks, triggering a scale that then weighs the employee before granting access to another locked door. This is an example of.

- A. mantrap.
- B. a bollard
- C. geofencing
- D. RFID.

Answer: A

Explanation:

A mantrap is a security device that consists of a small space with two sets of interlocking doors, such that the first set of doors must close before the second one opens. A mantrap can be used to control access to a data center by verifying the identity and weight of the person entering. A bollard is a sturdy post that prevents vehicles from entering a restricted area. Geofencing is a technology that uses GPS or RFID to create a virtual boundary around a location and trigger an action when a device crosses it. RFID is a technology that uses radio waves to identify and track objects or people. References:

? <https://www.techopedia.com/definition/16293/mantrap>

? <https://www.techopedia.com/definition/1437/bollard>

? <https://www.techopedia.com/definition/23961/geofencing>

? <https://www.techopedia.com/definition/506/radio-frequency-identification-rfid>

NEW QUESTION 4

A server administrator is connecting a new storage array to a server. The administrator has obtained multiple IP addresses for the array. Which of the following connection types is the server most likely using to connect to the array?

- A. eSATA
- B. USB
- C. FC
- D. iSCSI

Answer: D

Explanation:

iSCSI is a protocol that allows SCSI commands to be transmitted over IP networks, enabling remote access to storage devices. iSCSI uses IP addresses to identify and communicate with the storage array, so having multiple IP addresses for the array indicates that iSCSI is being used. eSATA, USB, and FC are other types of connections that use different protocols and connectors than iSCSI. References: CompTIA Server+ Certification Exam Objectives, Domain 3.0: Storage, Objective 3.1: Given a scenario, install and deploy primary storage devices based on given specifications and interfaces.

NEW QUESTION 5

A server administrator needs to deploy five VMs, all of which must have the same type of configuration. Which of the following would be the MOST efficient way to perform this task?

- A. Snapshot a VM.
- B. Use a physical host.
- C. Perform a P2V conversion.
- D. Use a VM template.

Answer: D

Explanation:

Deploying a virtual machine from a template creates a virtual machine that is a copy of the template. The new virtual machine has the virtual hardware, installed software, and other properties that are configured for the template.

Reference: https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vm_admin.doc/GUID-8254CD05-CC06-491D-BA56-A773A32A8130.html

The most efficient way to perform the task of deploying five VMs with the same type of configuration is to use a VM template. A template is a preconfigured virtual machine image that contains an operating system, applications, settings, and other components. A template can be used to create multiple identical or customized VMs quickly and easily, without having to install and configure each VM from scratch. A template can save time and ensure consistency across VMs.

NEW QUESTION 6

A company is implementing a check-in desk to heighten physical security. Which of the following access controls would be the most appropriate to facilitate this implementation?

- A. Security guards
- B. Security cameras
- C. Bollards
- D. An access control vestibule

Answer: D

Explanation:

An access control vestibule, or mantrap, is a type of physical access control that provides a space between two sets of interlocking doors. It is designed to prevent unauthorized individuals from following authorized individuals into facilities with controlled access, such as a check-in desk. The vestibule can be configured to limit the number of individuals who enter the controlled area and to verify their authorization for physical access. The other options are incorrect because they are not as effective as an access control vestibule in

facilitating the implementation of a check-in desk. Security guards, security cameras, and bollards are useful for monitoring, deterring, or preventing unauthorized access, but they do not provide the same level of control and verification as an access control vestibule

NEW QUESTION 7

DRAG DROP

A recent power Outage caused email services to go down. A server administrator also received alerts from the datacenter's UPS. After some investigation, the server administrator learned that each PDU was rated at a maximum Of 12A.

INSTRUCTIONS

Ensure power redundancy is implemented throughout each rack and UPS alarms are resolved. Ensure the maximum potential PDU consumption does not exceed 80% or 9.6A).

- * a. PDU selections must be changed using the pencil icon.
- * b. VM Hosts 1 and 2 and Mail Relay can be moved between racks.
- * c. Certain devices contain additional details

Data Center Racks 1 and 2

Show Question Reset All Answers

Rack 1

PDU A: 13A MAX PDU B: 6A MAX

Rack Switch 1

VM Host 1

VM Host 2

Mail Relay

SAN

PDU A

Rack 2

PDU A: 5A MAX PDU B: 8A MAX

Rack Switch 2

Domain Controller

?

?

PDU A

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Data Center Racks 1 and 2



NEW QUESTION 8

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 9

Which of the following testing exercises for disaster recovery is primarily used to discuss incident response strategies for critical systems without affecting production data?

- A. Tabletop
- B. Backup recovery test
- C. Lrverail over
- D. Hot-site visit audit

Answer: A

Explanation:

A tabletop exercise is a type of disaster recovery testing exercise that is primarily used to discuss incident response strategies for critical systems without affecting production data. A tabletop exercise is a discussion-based session where team members meet in an informal, classroom setting to review their roles and responsibilities during an emergency and their responses to a hypothetical scenario. A facilitator guides the participants through the discussion and evaluates the strengths and weaknesses of the preparedness program. A tabletop exercise does not involve any actual deployment of resources or activation of systems. A backup recovery test (B) is a type of disaster recovery testing exercise that involves restoring data from backup media to verify its integrity and availability. A backup recovery test may affect production data if it is not performed on a separate environment. A live failover is a type of disaster recovery testing exercise that involves switching operations from a primary site to a secondary site in case of a failure or disruption. A live failover may affect production data if it is not performed on a simulated environment. A hot-site visit audit (D) is a type of disaster recovery testing exercise that involves inspecting and evaluating a hot site, which is a backup location that has fully operational equipment and resources to resume business operations in case of a disaster. A hot-site visit audit does not involve any discussion of incident response strategies or simulation of scenarios. References: 1

<https://www.ready.gov/testing-exercises2> <https://www.ready.gov/exercises>

NEW QUESTION 10

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

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

Answer: C

Explanation:

The administrator should use 172.16.19.241 as an IP address to be in compliance with RFC 1918 standard. RFC 1918 defines three ranges of IP addresses that

are reserved for private internets, meaning they are not globally routable on the public Internet and can be used within an enterprise without any risk of conflict or overlap with other networks. These ranges are:

* 10.0.0.0 - 10.255.255.255 (10/8 prefix) 172.16.0.0 - 172.31.255.255 (172.16/12 prefix)

* 192.168.0.0 - 192.168.255.255 (192.168/16 prefix)

Out of these ranges, only 172.16.19.241 falls within one of them (172.16/12 prefix). The other options are either public IP addresses that belong to other organizations or networks (11.251.196.241, 171.245.198.241) or invalid IP addresses that do not conform to any standard (193.168.145.241).

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

NEW QUESTION 10

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

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

Answer: D

Explanation:

The server is not joined to the domain is causing the issue. A domain is a logical grouping of computers that share a common directory database and security policy on a network. Active Directory is a Microsoft technology that provides domain services for Windows-based computers. To use Active Directory credentials to log on to a server, the server must be joined to the domain that hosts Active Directory. If the server is not joined to the domain, it will not be able to authenticate with Active Directory and will only accept local accounts for logon. To join a server to a domain, the administrator must have a valid domain account with sufficient privileges and must know the name of the domain controller that hosts Active Directory.

NEW QUESTION 14

Which of the following backup types resets the archive bit each time it is run?

- A. Differential
- B. Snapshot
- C. Incremental
- D. Synthetic full

Answer: C

Explanation:

Incremental backup is a type of backup that only backs up the files that have changed since the last backup, whether it was a full or an incremental backup. Incremental backup resets the archive bit each time it is run, which means it clears the flag that indicates whether or not the file has been backed up. Incremental backup can save time and space compared to full backup, but it requires more time and resources to restore data from multiple backups. References: <https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-objectives> (Objective 3.1)

NEW QUESTION 18

A server administrator needs to configure a server on a network that will have no more than 30 available IP addresses. Which of the following subnet addresses will be the MOST efficient for this network?

- A. 255.255.255.0
- B. 255.255.255.128
- C. 255.255.255.224
- D. 255.255.255.252

Answer: C

Explanation:

The most efficient subnet address for a network that will have no more than 30 available IP addresses is 255.255.255.224. This subnet mask corresponds to a /27 prefix length, which means that 27 bits are used for the network portion and 5 bits are used for the host portion of an IP address. With 5 bits for hosts, there are $2^5 - 2 = 30$ possible host addresses per subnet, which meets the requirement. The other options are either too large or too small for the network size. Reference: <https://www.ibm.com/cloud/learn/subnet-mask>

NEW QUESTION 21

Which of the following concepts refers to prioritizing a connection that had previously worked successfully?

- A. Round robin
- B. SCP
- C. MRU
- D. Link aggregation

Answer: C

Explanation:

MRU, or Most Recently Used, is a concept that refers to prioritizing a connection that had previously worked successfully. It is often used in load balancing algorithms to distribute the workload among multiple servers or paths. MRU assumes that the most recently used connection is the most likely to be available and efficient, and therefore assigns the next request to that connection. This can help reduce latency and improve performance¹². The other options are incorrect because they do not refer to prioritizing a previous connection. Round robin is a concept that refers to distributing the workload equally among all available connections in a circular order¹². SCP, or Secure Copy Protocol, is a concept that refers to transferring files securely between hosts using encryption³. Link aggregation is a concept that refers to combining multiple physical links into a single logical link to increase bandwidth and redundancy⁴.

NEW QUESTION 25

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 30

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

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

Answer: BF

Explanation:

The server has a CMOS battery failure and the time synchronization service is disabled on the servers. The CMOS battery is a small battery on the motherboard that powers the BIOS settings and keeps track of the date and time when the server is powered off. If the CMOS battery fails, the server will lose its configuration and display an incorrect date and time when it is powered on. This can cause access issues for users and applications that rely on accurate time stamps. The time synchronization service is a service that synchronizes the system clock with a reliable external time source, such as a network time protocol (NTP) server. If the time synchronization service is disabled on the servers, they will not be able to update their clocks automatically and may drift out of sync with each other and with the network. This can also cause access issues for users and applications that require consistent and accurate time across the network.

NEW QUESTION 33

A systems administrator notices a newly added server cannot see any of the LUNs on the SAN. The SAN switch and the local HBA do not display any link lights. Which of the following is most likely the issue?

- A. A single-mode fiber cable is used in place of multimode.
- B. The switchport is on the wrong virtual SAN.
- C. The HBA driver needs to be installed on the server.
- D. The zoning on the fiber switch is wrong.

Answer: A

Explanation:

The most likely issue that prevents the newly added server from seeing any of the LUNs on the SAN is that a single-mode fiber cable is used in place of multimode. A single-mode fiber cable is a type of optical fiber cable that has a small core diameter and allows only one mode of light to propagate through it. A single-mode fiber cable can transmit data over long distances at high speeds, but it requires more expensive transceivers and connectors than multimode fiber cables. A multimode fiber cable is a type of optical fiber cable that has a larger core diameter and allows multiple modes of light to propagate through it. A multimode fiber cable can transmit data over short distances at lower speeds than single-mode fiber cables, but it is more compatible and cost-effective than single-mode fiber cables. If a single-mode fiber cable is used in place of multimode, it can cause signal loss, attenuation, or mismatch between the devices. References: [CompTIA Server+ Certification Exam Objectives], Domain 3.0: Storage, Objective 3.2: Given a scenario, compare and contrast various storage technologies.

NEW QUESTION 35

A server room with many racks of servers is managed remotely with occasional on-site support. Which of the following would be the MOST cost-effective option to administer and troubleshoot network problems locally on the servers?

- A. Management port
- B. Crash cart
- C. IP KVM
- D. KVM

Answer: C

Explanation:

An IP KVM (keyboard, video, mouse) is a device that allows remote access and control of multiple servers over a network using a web browser or a client software. An IP KVM is a cost-effective option to administer and troubleshoot network problems locally on the servers, as it eliminates the need for physical presence or dedicated hardware for each server. A management port (A) is a network interface that is used for out-of-band management of network devices, such as routers or switches. A management port does not provide local access to servers. A crash cart (B) is a mobile unit that contains a monitor, keyboard, mouse, and other tools for troubleshooting servers in a data center. A crash cart requires physical access to each server and may not be cost-effective for many racks of servers. A KVM (D) is a device that allows switching between multiple servers using a single keyboard, video, and mouse. A KVM does not provide remote access over a network and requires physical connection to each server. References: <https://www.enterprisestorageforum.com/management/best-data-storage-solutions-and-software-2021/><https://www.microsoft.com/en-us/microsoft-365/business-insights-ideas/resources/cloud-storage-vs-on-premises-servers>

NEW QUESTION 38

A server room contains ten physical servers that are running applications and a cluster of three dedicated hypervisors. The hypervisors are new and only have 10% utilization. The Chief Financial Officer has asked that the IT department do what it can to cut back on power consumption and maintenance costs in the data center. Which of the following would address the request with minimal server downtime?

- A. Unplug the power cables from the redundant power supplies, leaving just the minimum required.
- B. Convert the physical servers to the hypervisors and retire the ten servers.
- C. Reimage the physical servers and retire all ten servers after the migration is complete.
- D. Convert the ten servers to power-efficient core editions.

Answer: B

Explanation:

This option would reduce power consumption and maintenance costs by consolidating the physical servers into virtual machines on the hypervisors. This would also free up space and resources in the data center. The other options would either not address the request, increase power consumption, or require more maintenance.

NEW QUESTION 42

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

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

Answer: B

Explanation:

Cloning the original VM is the most efficient way to provide a VM with high availability. Cloning is the process of creating an exact copy of a VM, including its configuration, operating system, applications, and data. A cloned VM can be used as a backup or a replica of the original VM, and can be powered on and run independently. Cloning can be done quickly and easily using vSphere tools or other third-party software. By cloning the original VM and placing it on a different host server or availability zone, the technician can ensure that if the original VM fails, the cloned VM can take over its role and provide uninterrupted service to the users and applications.

NEW QUESTION 43

A server is only able to connect to a gigabit switch at 100Mb. Other devices are able to access the network port at full gigabit speeds, and when the server is brought to another location, it is able to connect at full gigabit speed. Which of the following should an administrator check first?

- A. The switch management
- B. The VLAN configuration
- C. The network cable
- D. The network drivers

Answer: C

Explanation:

The first thing that the administrator should check is the network cable. The network cable is a physical medium that connects a server to a switch or other network device. The network cable can affect the speed and quality of the network connection, depending on its type, length, and condition. If the network cable is damaged, faulty, or incompatible, it can cause the server to connect at a lower speed than expected. Therefore, the administrator should check the network cable for any signs of wear, tear, or mismatch, and replace it if necessary.

References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 2, Lesson 2.1, Objective 2.1

NEW QUESTION 47

A security administrator ran a port scanning tool against a virtual server that is hosting a secure website. A list of open ports was provided as documentation. The management team has requested that non-essential ports be disabled on the firewall. Which of the following ports must remain open?

- A. 25
- B. 443
- C. 3389
- D. 8080

Answer: B

Explanation:

The port that must remain open for a secure website is port 443. Port 443 is used by Hypertext Transfer Protocol Secure (HTTPS), which is an extension of HTTP that encrypts and authenticates the communication between a web server and a web browser. HTTPS ensures that the data transmitted over the web is protected from eavesdropping, tampering, or spoofing. Therefore, port 443 must remain open for a secure website to function properly.

References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 2, Lesson 2.2, Objective 2.2

NEW QUESTION 52

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 56

A server administrator is creating a new server that will be used to house customer sales records. Which of the following roles will MOST likely be installed on the server?

- A. Print
- B. File
- C. Database
- D. Messaging

Answer: C

Explanation:

A database server is a server that hosts a database management system (DBMS) that stores, organizes, and manipulates data. A database server is suitable for housing customer sales records, as it can provide fast and secure access, query and analysis capabilities, backup and recovery options, and scalability and performance optimization. Some examples of database servers are Microsoft SQL Server, Oracle Database, MySQL, and PostgreSQL. Verified References: [What is a Database Server?]

NEW QUESTION 57

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
```

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 62

A server administrator is deploying a new server that has two hard drives on which to install the OS. Which of the following RAID configurations should be used to provide redundancy for the OS?

- A. RAID 0
- B. RAID 1
- C. RAID 5
- D. RAID 6

Answer: B

Explanation:

RAID 1 (mirroring) is a RAID configuration that should be used to provide redundancy for the OS on a server that has two hard drives on which to install the OS. RAID 1 (mirroring) is a configuration that duplicates data across two or more drives. It provides fault tolerance and improves read performance, but reduces storage capacity by half. If one drive fails in RAID 1, the other drive can continue to operate without data loss or system downtime. RAID 0 (striping) is a configuration that splits data across two or more drives without parity or redundancy. It improves performance but offers no fault tolerance. If one drive fails in RAID 0, all data is lost and the system cannot boot. RAID 5 (striping with parity) is a configuration that stripes data across three or more drives with parity information. It provides fault tolerance and improves performance, but reduces storage capacity by one drive's worth of space. RAID 5 can tolerate one drive failure without data loss, but not two or more. RAID 6 (striping with double parity) is a configuration that stripes data across four or more drives with double parity information. It provides fault tolerance and improves performance, but reduces storage capacity by two drives' worth of space. RAID 6 can tolerate two drive failures without data loss, but not three or more. References: <https://www.howtogeek.com/199068/how-to-upgrade-your-existing-hard-drive-in-under-an-hour/>

NEW QUESTION 63

Which of the following commands would MOST likely be used to register a new service on a Windows OS?

- A. set-service
- B. net
- C. sc
- D. services.msc

Answer: C

Explanation:

The sc command is used to create, delete, start, stop, pause, or query services on a Windows OS. It can also be used to register a new service by using the create option. References: <https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/sc-create>

NEW QUESTION 64

A server administrator added a new drive to a server. However, the drive is not showing up as available. Which of the following does the administrator need to do to make the drive available?

- A. Partition the drive.
- B. Create a new disk quota.
- C. Configure the drive as dynamic.
- D. Set the compression.

Answer: A

Explanation:

To make a new drive available on a server, the administrator needs to partition the drive first. Partitioning is a process that divides the drive into one or more logical sections that can be formatted and assigned drive letters or mount points. Partitioning can be done using tools such as Disk Management on Windows or fdisk on Linux. Creating a new disk quota would not help, as disk quotas are used to limit the amount of disk space that users or groups can use on a partition. Configuring the drive as dynamic would not help either, as dynamic disks are used to create volumes that span multiple disks or use RAID features. Setting the compression would not help, as compression is used to reduce the size of files on a partition. References: <https://www.howtogeek.com/school/using-windows-admin-tools-like-a-pro/lesson2/> <https://www.howtogeek.com/howto/17001/how-to-format-a-usb-drive-in-ubuntu-using-gparted/>

NEW QUESTION 69

A technician is attempting to log in to a Linux server as root but cannot remember the administrator password. Which of the following is the LEAST destructive method of resetting the administrator password?

- A. Boot using a Linux live CD and mount the hard disk to /mn
- B. Change to the /mnt/etc directory
- C. Edit the passwd file found in that directory.
- D. Reinstall the OS in overlay mod
- E. Reset the root password from the install GUI screen.
- F. Adjust the GRUB boot parameters to boot into single-user mod
- G. Run passwd from the command prompt.
- H. Boot using a Linux live CD and mount the hard disk to /mn
- I. SCP the /etc directory from a known accessible server to /mnt/etc.

Answer: C

Explanation:

This is the least destructive method of resetting the administrator password because it does not require modifying any files or reinstalling the OS. It only requires changing the boot parameters temporarily and running a command to change the password. References: https://wiki.archlinux.org/title/Reset_lost_root_password#Using_GRUB

NEW QUESTION 72

Users cannot access a new server by name, but the server does respond to a ping request using its IP address. All the user workstations receive their IP information from a DHCP server. Which of the following would be the best step to perform NEXT?

- A. Run the tracert command from a workstation.
- B. Examine the DNS to see if the new server record exists.
- C. Correct the missing DHCP scope.
- D. Update the workstation hosts file.

Answer: B

Explanation:

If users cannot access a new server by name, but the server does respond to a ping request using its IP address, it means that there is a problem with name

resolution. The DNS (Domain Name System) is a service that maps hostnames to IP addresses and vice versa. Therefore, the best step to perform next is to examine the DNS to see if the new server record exists and matches its IP address. If not, the DNS record needs to be added or updated accordingly. Running the tracert command from a workstation would not help with name resolution, as it only shows the route taken by packets to reach a destination by IP address. Correcting the missing DHCP scope would not help either, as DHCP (Dynamic Host Configuration Protocol) only assigns IP addresses and other network settings to clients, but does not resolve names. Updating the workstation hosts file would be a temporary workaround, but not a permanent solution, as it would require manually editing every workstation's hosts file with the new server's name and IP address. References: <https://www.howtogeek.com/164981/how-to-use-nslookup-to-check-domain-name-information-in-microsoft-windows/><https://www.howtogeek.com/howto/27350/beginner-geek-how-to-edit-your-hosts-file/>

NEW QUESTION 75

Which of the following license types most commonly describes a product that incurs a yearly cost regardless of how much it is used?

- A. Physical
- B. Subscription
- C. Open-source
- D. Per instance
- E. Per concurrent user

Answer: B

Explanation:

A subscription license is a type of license that grants the user the right to use a product or service for a fixed period of time, usually a year. The user pays a recurring fee, regardless of how much they use the product or service. Subscription licenses are common for cloud-based software and services, such as Microsoft 365 or DocuSign2.

References = 1: Compare All Microsoft 365 Plans (Formerly Office 365) - Microsoft Store(<https://www.microsoft.com/en-us/microsoft-365/buy/compare-all-microsoft-365-products>) 2: DocuSign Pricing | eSignature Plans for Personal & Business(<https://ecom.docusign.com/plans-and-pricing/esignature>)

NEW QUESTION 76

A datacenter technician is attempting to troubleshoot a server that keeps crashing. The server runs normally for approximately five minutes, but then it crashes. After restoring the server to operation, the same cycle repeats. The technician confirms none of the configurations have changed, and the load on the server is steady from power-on until the crash. Which of the following will MOST likely resolve the issue?

- A. Reseating any expansion cards in the server
- B. Replacing the failing hard drive
- C. Reinstalling the heat sink with new thermal paste
- D. Restoring the server from the latest full backup

Answer: C

Explanation:

The most likely solution to resolve the issue of the server crashing after running normally for approximately five minutes is to reinstall the heat sink with new thermal paste. A heat sink is a device that dissipates heat from a component, such as a processor or a graphics card, by transferring it to a cooling medium, such as air or liquid. A heat sink is usually attached to the component using thermal paste, which is a substance that fills the gaps between the heat sink and the component and improves thermal conductivity. Thermal paste can degrade over time and lose its effectiveness, resulting in overheating and performance issues. If a server crashes after running for a short period of time, it may indicate that the processor is overheating due to insufficient cooling. To resolve this issue, the technician should remove the heat sink, clean the old thermal paste, apply new thermal paste, and reinstall the heat sink.

NEW QUESTION 77

A storage administrator needs to implement SAN-based shared storage that can transmit at 16Gb over an optical connection. Which of the following connectivity options would BEST meet this requirement?

- A. Fibre Channel
- B. FCoE
- C. iSCSI
- D. eSATA

Answer: A

Explanation:

Fibre Channel is a connectivity option that can transmit at 16Gb over an optical connection for SAN-based shared storage. Fibre Channel is a high-speed network technology that provides reliable and secure data transfer between servers and storage devices. Fibre Channel uses optical fiber cables to connect devices and supports various topologies and protocols. FCoE is another connectivity option that uses Fibre Channel over Ethernet, which encapsulates Fibre Channel frames into Ethernet packets. FCoE can also transmit at 16Gb over an optical connection, but it requires a converged network adapter (CNA) and a lossless Ethernet network. iSCSI is another connectivity option that uses SCSI commands over IP networks, which can use either copper or optical cables. iSCSI can transmit at 10Gb or 40Gb over an optical connection, but it has higher latency and lower performance than Fibre Channel. eSATA is another connectivity option that uses SATA commands over external cables, which are usually copper. eSATA can transmit at 6Gb over a copper connection, but it has limited cable length and device support compared to Fibre Channel. References:

? <https://www.ibm.com/topics/storage-area-network>

? <https://www.techopedia.com/definition/1369/fibre-channel-fc>

? <https://www.techopedia.com/definition/1368/fibre-channel-over-ethernet-fcoe>

? <https://www.techopedia.com/definition/1367/internet-small-computer-system-interface-iscsi>

? <https://www.techopedia.com/definition/1366/external-serial-advanced-technology-attachment-esata>

NEW QUESTION 78

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 79

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 demand¹².

? Scalability: Cloud IaaS can offer flexible and elastic scalability of resources, as

users can easily provision or deprovision resources according to their needs and workload¹².

? Availability: Cloud IaaS can ensure high availability and reliability of the

application, as users can leverage the cloud provider's redundant and geographically distributed infrastructure¹².

? 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 operations¹².

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 83

A technician is tasked with upgrading 24 hosts simultaneously with a Type 1 hypervisor. Which of the following protocols should the technician use for this upgrade?

- A. VPN
- B. TFTP
- C. SSH
- D. HTTP

Answer: B

Explanation:

TFTP (Trivial File Transfer Protocol) is a simple and lightweight protocol that can be used to transfer files over a network. TFTP is often used to upgrade firmware or software on network devices, such as routers, switches, or servers. TFTP can also be used to install a Type 1 hypervisor, such as VMware ESXi, on multiple hosts simultaneously¹². References = 1: How to Install VMware ESXi Type 1 Hypervisor - MatthewEaton.net(<https://mattheweaton.net/posts/how-to-install-vmware-esxi-type-1-hypervisor/>) 2: Explore Type 1 Hypervisors - Set Up Virtual Machines Using VirtualBox and vSphere -

OpenClassrooms(<https://openclassrooms.com/en/courses/7163136-set-up-virtual-machines-using-virtualbox-and-vsphere/7358546-explore-type-1-hypervisors>)

NEW QUESTION 84

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

NEW QUESTION 88

A server administrator just installed a new physical server and needs to harden the applications on the server. Which of the following best describes a method of application hardening?

- A. Install the latest patches.
- B. Disable unneeded hardware.
- C. Set the boot order.
- D. Enable a BIOS password.

Answer: A

Explanation:

A method of application hardening is installing the latest patches. Application hardening is a process of reducing the attack surface and vulnerabilities of an application by applying security measures and best practices. Installing the latest patches is one way to harden an application, as patches are updates that fix bugs, errors, or security issues in an application. By installing the latest patches, an application can be protected from known exploits or threats.

References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 3, Lesson 3.5, Objective 3.5

NEW QUESTION 89

A systems administrator is setting up a server farm for a new company. The company has a public range of IP addresses and uses the addresses internally. Which of the following IP addresses best fits this scenario?

- A. 10.3.7.27
- B. 127.0.0.1
- C. 192.168.7.1
- D. 216.176.128.10

Answer: D

Explanation:

The IP address that best fits this scenario is 216.176.128.10. This is a public IP address that belongs to a range of addresses that are assigned and registered by an Internet service provider (ISP) and can be accessed from anywhere on the Internet. The company has a public range of IP addresses and uses them internally, which means that they do not use private IP addresses or network address translation (NAT) to communicate within their network.

References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 2, Lesson 2.2, Objective 2.2

NEW QUESTION 90

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

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

Answer: B

Explanation:

Shell is the most appropriate scripting language to use for a logon script for a Linux box. Shell is a generic term for a command-line interpreter that allows users to interact with the operating system by typing commands and executing scripts. Shell scripts are files that contain a series of commands and instructions that can be executed by a shell. Shell scripts are commonly used for automating tasks, such as logon scripts that run when a user logs on to a system. There are different types of shells available for Linux systems, such as Bash, Ksh, Zsh, etc., but they all share a similar syntax and functionality.

NEW QUESTION 91

A data center environment currently hosts more than 100 servers that include homegrown and commercial software. The management team has asked the server administrator to find a way to eliminate all company-owned data centers. Which of the following models will the administrator most likely choose to meet this need?

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

Answer: C

Explanation:

A public cloud model will most likely meet the need of eliminating all company-owned data centers. A public cloud is a type of cloud computing service that is provided by a third-party vendor over the internet. A public cloud offers scalability, flexibility, and cost-effectiveness for hosting servers and applications, as the customers only pay for the resources they use and do not have to maintain their own infrastructure. A public cloud can also provide high availability, security, and performance for the servers and applications, as the vendor manages the underlying hardware and software. A public cloud can support various types of services, such as software as a service (SaaS), platform as a service (PaaS), or infrastructure as a service (IaaS). References: [CompTIA Server+ Certification Exam Objectives], Domain 1.0: Server Administration, Objective 1.2: Given a scenario, compare and contrast server roles and requirements for each.

NEW QUESTION 95

A company uses a hot-site, disaster-recovery model. Which of the following types of data replication is required?

- A. Asynchronous
- B. Incremental
- C. Application consistent
- D. Constant

Answer: D

Explanation:

The type of data replication that is required for a hot-site disaster recovery model is constant. A hot site is a type of disaster recovery site that has fully operational IT infrastructure and equipment that can take over the primary site's functions immediately in case of a disaster or disruption. A hot site requires constant data replication between the primary site and the hot site to ensure that the data is up-to-date and consistent. Constant data replication means that any changes made to the data at the primary site are immediately copied to the hot site without any delay or lag.

NEW QUESTION 100

An analyst is planning a new point-to-point Ethernet connection between a university campus and a newly acquired space downtown that is about 5mi (8km) away. The analyst would like the fastest possible connection speed. Which of the following would best meet the analyst's needs?

- A. 1000BASE-LX 1Gb single-mode plenum fiber connection
- B. 10GBASE-T 10Gb copper plenum Ethernet connection
- C. 1000BASE-T 1Gb copper non-plenum Ethernet connection
- D. 10GBASE-SR 10Gb multimode plenum fiber connection

Answer: A

Explanation:

A 1000BASE-LX 1Gb single-mode plenum fiber connection would best meet the analyst's needs for a new point-to-point Ethernet connection between a university campus and a newly acquired space downtown that is about 5mi (8km) away. A 1000BASE-LX is a type of Ethernet standard that supports data transmission at 1 gigabit per second over single-mode fiber cables using long wavelength lasers. A single-mode fiber cable is a type of optical fiber cable that has a small core diameter and allows only one mode of light to propagate through it. A single-mode fiber cable can transmit data over long distances at high speeds, but it requires more expensive transceivers and connectors than multimode fiber cables. A plenum fiber cable is a type of optical fiber cable that has a special coating that prevents the spread of fire or toxic fumes in case of burning. A plenum fiber cable is suitable for installation in plenum spaces, which are areas used for air circulation in buildings, such as above ceilings or below floors. References: [CompTIA Server+ Certification Exam Objectives], Domain 4.0: Networking, Objective 4.2: Given a scenario involving server networking issues (e.g., network interface card failure), troubleshoot using appropriate tools.

NEW QUESTION 102

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

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

Answer: B

Explanation:

The administrator needs to install SSL to make the website trusted. SSL stands for Secure Sockets Layer, which is an encryption-based Internet security protocol that ensures privacy, authentication, and data integrity in web communications. SSL enables HTTPS (Hypertext Transfer Protocol Secure), which is a secure version of HTTP (Hypertext Transfer Protocol) that encrypts the data exchanged between a web browser and a web server. SSL also uses digital certificates to verify the identity of the web server and establish trust with the web browser. A web server that implements SSL has HTTPS in its URL instead of HTTP and displays a padlock icon or a green bar in the browser's address bar.

NEW QUESTION 106

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 107

An administrator has been asked to deploy a database server that provides the highest performance with fault tolerance. Which of the following RAID levels will fulfill this request?

- A. RAID0
- B. RAID1
- C. RAID5

- D. RAID 6
- E. RAID 10

Answer: E

Explanation:

RAID 10 is the best option to deploy a database server that provides the highest performance with fault tolerance. RAID 10 is a type of RAID level that combines RAID 1 (mirroring) and RAID 0 (striping) to create an array of mirrored stripes. RAID 10 offers high performance by distributing data across multiple disks in parallel (striping), which improves read/write speed and I/O operations. RAID 10 also offers fault tolerance by duplicating data across two or more disks in each stripe (mirroring), which provides redundancy and data protection in case of disk failure. RAID 10 requires at least four disks to implement and has a high storage overhead, as half of the disk space is used for mirroring. References: [CompTIA Server+ Certification Exam Objectives]

NEW QUESTION 110

A server administrator is creating a script that will move files only if they were created before a date input by the user. Which of the following constructs will allow the script to apply this test until all available files are assessed?

- A. Variable
- B. Loop
- C. Comparator
- D. Conditional

Answer: B

Explanation:

A loop is a script construct that allows the script to repeat a block of code until a certain condition is met or for a specified number of times. A loop can be used to apply a test to each file in a directory and move the files that meet the criteria. For example, in a bash script, a loop can be written as:

```
#!/bin/bash
# Ask the user for the date echo"Enter the date (YYYY-MM-DD):" readdate
# Loop through all the files in the current directory forfilein*
do
# Check if the file was created before the date if[[ $(date-r"$file"+%F) <$date]]
then
# Move the file to another location mv"$file"/path/to/destination
fi done Copy
```

A variable is a script construct that allows the script to store and manipulate data. A variable can be used to store the date input by the user, but it cannot apply a test to each file¹

A comparator is a script construct that allows the script to compare two values and determine their relationship. A comparator can be used to check if a file was created before

the date, but it cannot repeat the test for all files¹

A conditional is a script construct that allows the script to execute different blocks of code based on certain conditions. A conditional can be used to decide whether to move a file or not, but it cannot iterate over all files¹

1: CompTIA Server+ Certification Exam Objectives

NEW QUESTION 114

A server administrator just installed a new physical server and needs to harden the OS. Which of the following best describes the OS hardening method?

- A. Apply security updates.
- B. Disable unneeded hardware.
- C. Set a BIOS password.
- D. Configure the boot order.

Answer: A

Explanation:

Applying security updates is one of the common operating system hardening methods that can help protect the OS from cyberattacks and vulnerabilities. Security updates are released by the OS developer to fix bugs, patch security holes, and improve performance. By installing the latest updates, the server administrator can ensure that the OS is up to date and secure¹².

NEW QUESTION 115

A technician has moved a data drive from a new Windows server to an older Windows server. The hardware recognizes the drive, but the data is not visible to the OS. Which of the following is the most likely cause of the issue?

- A. The disk uses GPT.
- B. The partition is formatted with ext4.
- C. The partition is formatted with FAT32.
- D. The disk uses MBR.

Answer: A

Explanation:

The most likely cause of the issue is that the disk uses GPT. GPT stands for GUID Partition Table, which is a newer standard for disk partitioning that supports larger disks and more partitions than the older MBR (Master Boot Record) standard¹. However, GPT is not compatible with some older operating systems, such as Windows XP or Windows Server 2003². Therefore, if the data drive was formatted with GPT on a new Windows server and then moved to an older Windows server, the older server may not be able to recognize the GPT partitions and access the data on the drive.

The partition being formatted with ext4, FAT32, or MBR are not likely causes of the issue. Ext4 is a file system that is commonly used on Linux-based systems, but it can also be read by Windows with some third-party software³. FAT32 is a file system that is widely compatible with most operating systems and devices, but it has some limitations such as a maximum file size of 4 GB and a maximum partition size of 8 TB⁴. MBR is not a file system, but a partitioning scheme that can support various file systems such as NTFS, FAT32, or exFAT⁵. However, MBR has some disadvantages compared to GPT, such as a maximum disk size of 2 TB and a maximum number of primary partitions of four¹.

NEW QUESTION 119

Which of the following BEST measures how much downtime an organization can tolerate during an unplanned outage?

- A. SLA
- B. BIA
- C. RTO
- D. MTTR

Answer: C

Explanation:

RTO (Recovery Time Objective) is a measure of how much downtime an organization can tolerate during an unplanned outage. It is the maximum time allowed for restoring normal operations after a disaster. RTO is one of the key metrics for disaster recovery planning and testing. SLA (Service Level Agreement) is a contract that defines the expected level of service and performance between a provider and a customer. BIA (Business Impact Analysis) is a process that identifies and evaluates the potential effects of a disaster on critical business functions and processes. MTTR (Mean Time To Repair) is a measure of how long it takes to fix a failed component or system. References: <https://parachute.cloud/rto-vs-rpo/> <https://www.techopedia.com/definition/13622/service-level-agreement-sla> <https://www.techopedia.com/definition/1032/business-impact-analysis-bia> <https://www.techopedia.com/definition/8239/mean-time-to-repair-mttr>

NEW QUESTION 121

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 125

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

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

Answer: D

Explanation:

The administrator should use SCP to transfer log files from a Linux server to a Windows workstation. SCP (Secure Copy Protocol) is a protocol that allows secure file transfer between two devices using SSH (Secure Shell) encryption. SCP can transfer files between different operating systems, such as Linux and Windows, as long as both devices have an SSH client installed. SCP can also preserve file attributes, such as permissions and timestamps, during the transfer.

NEW QUESTION 126

A server administrator needs to keep a copy of an important fileshare that can be used to restore the share as quickly as possible. Which of the following is the BEST solution?

- A. Copy the fileshare to an LTO-4 tape drive
- B. Configure a new incremental backup job for the fileshare
- C. Create an additional partition and move a copy of the fileshare
- D. Create a snapshot of the fileshare

Answer: D

Explanation:

The best solution to keep a copy of an important fileshare that can be used to restore the share as quickly as possible is to create a snapshot of the fileshare. A snapshot is a point-in-time copy of a file system or a volume that captures the state and data of the fileshare at a specific moment. A snapshot can be created instantly and with minimal overhead, as it only stores the changes made to the fileshare after the snapshot was taken. A snapshot can be used to restore the fileshare to its previous state in case of data loss or corruption.

NEW QUESTION 130

Which of the following backup types copies changed data from a server and then combines the backups on the backup target?

- A. Differential
- B. Incremental
- C. Synthetic full
- D. Snapshot

Answer: C

Explanation:

A synthetic full backup is a type of backup that copies changed data from a server and then combines the backups on the backup target. This way, the backup target always has a full backup of the server, without requiring a full backup to be performed over the network. A synthetic full backup reduces the network bandwidth and time required for backups, while also simplifying the restoration process.

NEW QUESTION 133

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 `logrotate` directive from the `audit.conf` 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 134

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 137

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

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

Answer: A

Explanation:

The administrator should use an IP KVM to access the server's console remotely for bare-metal maintenance. An IP KVM stands for Internet Protocol Keyboard Video Mouse, which is a device that allows remote control of a server's keyboard, video, and mouse over a network connection, such as LAN or Internet. An IP KVM enables an administrator to perform tasks such as BIOS configuration, boot sequence selection, operating system installation, etc., without being physically present at the server location. The other options are not suitable for bare-metal maintenance because they require either physical access to the server (a crash cart) or an operating system running on the server (VNC, RDP, SSH). A crash cart is a mobile unit that contains a monitor, keyboard, mouse, and cables that can be plugged into a server for direct access to its console. VNC stands for Virtual Network Computing, which is a software that allows remote desktop sharing and control over a network connection using a graphical user interface (GUI). RDP stands for Remote Desktop Protocol, which is a protocol that allows remote desktop access and control over a network connection using a GUI or command-line interface (CLI). SSH stands for Secure Shell, which is a protocol that allows secure remote login and command execution over a network connection using a CLI.

NEW QUESTION 139

A server administrator is experiencing difficulty configuring MySQL on a Linux server. The administrator issues the `getenforce` command and receives the following output:

```
># Enforcing
```

Which of the following commands should the administrator issue to configure MySQL successfully?

- A. `setenforce 0`
- B. `setenforce permissive`
- C. `setenforce 1`
- D. `setenforce disabled`

Answer: A

Explanation:

The command that the administrator should issue to configure MySQL successfully is `setenforce 0`. This command sets the SELinux (Security-Enhanced Linux) mode to permissive, which means that SELinux will not enforce its security policies and will only log any violations. SELinux is a feature that provides mandatory

access control (MAC) for Linux systems, which can enhance the security and prevent unauthorized access or modification of files and processes. However, SELinux can also interfere with some applications or services that require specific permissions or ports that are not allowed by SELinux by default. In this case, MySQL may not be able to run properly due to SELinux restrictions. To resolve this issue, the administrator can either disable SELinux temporarily by using `setenforce 0`, or permanently by editing the `/etc/selinux/config` file and setting `SELINUX=disabled`. Alternatively, the administrator can configure SELinux to allow MySQL

to run by using commands such as `semanage` or `setsebool`.

Reference:

<https://blogs.oracle.com/mysql/selinux-and-mysql-v2>

NEW QUESTION 142

A company stores extremely sensitive data on an air-gapped system. Which of the following can be implemented to increase security against a potential insider threat?

- A. Two-person Integrity
- B. SSO
- C. SIEM
- D. Faraday cage
- E. MFA

Answer: A

Explanation:

Two-person integrity is a security measure that can be implemented to increase security against a potential insider threat on an air-gapped system. An air-gapped system is a system that is isolated from any network connection and can only be accessed physically. An insider threat is a malicious actor who has authorized access to an organization's system or data and uses it for unauthorized or harmful purposes. Two-person integrity is a system of storage and handling that requires the presence of at least two authorized persons, each capable of detecting incorrect or unauthorized security procedures, for accessing certain sensitive data or material. This way, no single person can compromise the security or integrity of the data or material without being noticed by another person. SSO (Single Sign-On) is a feature that allows users to access multiple applications or systems with one set of credentials, but it does not prevent insider threats. SIEM (Security Information and Event Management) is a tool that collects and analyzes log data from various sources to detect and respond to security incidents, but it does not work on air-gapped systems. A Faraday cage is a structure that blocks electromagnetic signals from entering or leaving, but it does not prevent physical access or insider threats. MFA (Multi-Factor Authentication) is a method that requires users to provide two or more pieces of evidence to verify their identity, such as something they know, something they have, or something they are, but it does not prevent insider threats. References: <https://www.howtogeek.com/169080/air-gap-how-to-isolate-a-computer-to-protect-it-from-hackers/> <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 147

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 148

A technician has moved a data drive from a new Windows server to an older Windows server. The hardware recognizes the drive, but the data is not visible to the OS. Which of the following is the MOST Likely cause of the issue?

- A. The disk uses GPT.
- B. The partition is formatted with ext4.
- C. The partition is formatted with FAT32.
- D. The disk uses MBR.

Answer: A

Explanation:

GPT (GUID Partition Table) is a partitioning scheme that allows creating partitions on large hard drives (more than 2 TB). It supports up to 128 partitions per drive and uses 64-bit addresses to locate them. However, GPT is not compatible with older versions of Windows, such as Windows XP or Windows Server 2003, which

use MBR (Master Boot Record) as the partitioning scheme. If a disk uses GPT, it may not be recognized or accessible by an older Windows server. Verified References: [GPT], [MBR]

NEW QUESTION 149

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 Objectives¹, page 8
- ? RAID Levels and Types Explained: Advantages and Disadvantages²
- ? RAID Levels & Fault Tolerance³

NEW QUESTION 153

A junior administrator needs to configure a single RAID 5 volume out of four 200GB drives attached to the server using the maximum possible capacity. Upon completion, the server reports that all drives were used, and the approximate volume size is 400GB. Which of the following BEST describes the result of this configuration?

- A. RAID 0 was configured by mistake.
- B. RAID 5 was configured properly.
- C. JBOD was configured by mistake.
- D. RAID 10 was configured by mistake.

Answer: B

Explanation:

The output of the configuration shows that RAID 5 was configured properly using four 200GB drives. The approximate volume size of 400GB is correct, since RAID 5 uses one disk for parity and the rest for data. Therefore, the usable storage capacity is three-fourths of the total capacity, which is 600GB out of 800GB. The other RAID levels given would result in different volume sizes: RAID 0 would result in 800GB, RAID 1 would result in 200GB, and JBOD would result in an error since it does not support multiple drives in a single volume. References: https://en.wikipedia.org/wiki/Standard_RAID_levels#RAID_5

NEW QUESTION 155

A technician is configuring a point-to-point heartbeat connection between two servers using IP addressing. Which of the following is the most efficient subnet mask for this connection?

- A. /28
- B. /29
- C. /30
- D. /32

Answer: C

Explanation:

The most efficient subnet mask for a point-to-point heartbeat connection between two servers using IP addressing is /30. A /30 subnet mask has 255.255.255.252 as its decimal representation and 11111111.11111111.11111111.11111100 as its binary representation. This means that there are only two bits available for the host portion of the IP address, which allows for four possible combinations: 00, 01, 10, and 11. However, the first and the last combinations are reserved for the network address and the broadcast address, respectively. Therefore, only two IP addresses are usable for the point-to-point connection, which is the minimum required for such a link. A /30 subnet mask is also known as a point-to-point prefix because it is commonly used for point-to-point links between routers or servers¹.

A /28 subnet mask has 255.255.255.240 as its decimal representation and 11111111.11111111.11111111.11110000 as its binary representation. This means that there are four bits available for the host portion of the IP address, which allows for 16 possible combinations. However, two of them are reserved for the network address and the broadcast address, respectively. Therefore, 14 IP addresses are usable for the subnet, which is more than needed for a point-to-point connection and would result in wasted addresses.

A /29 subnet mask has 255.255.255.248 as its decimal representation and 11111111.11111111.11111111.11111000 as its binary representation. This means that there are three bits available for the host portion of the IP address, which allows for eight possible combinations. However, two of them are reserved for the network address and the broadcast address, respectively. Therefore, six IP addresses are usable for the subnet, which is still more than needed for a point-to-point connection and would result in wasted addresses.

A /32 subnet mask has 255.255.255.255 as its decimal representation and 11111111.11111111.11111111.11111111 as its binary representation. This means that there are no bits available for the host portion of the IP address, which allows for only one possible combination: all ones. Therefore, only one IP address is usable for the subnet, which is not enough for a point-to-point connection and would result in an invalid configuration.

Therefore, a /30 subnet mask is the most efficient choice for a point-to-point heartbeat connection between two servers using IP addressing because it provides exactly two usable IP addresses without wasting any addresses or creating any conflicts¹.

NEW QUESTION 160

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 162

A technician is unable to access a server's package repository internally or externally. Which of the following are the MOST likely reasons? (Choose two.)

- A. The server has an architecture mismatch
- B. The system time is not synchronized
- C. The technician does not have sufficient privileges
- D. The external firewall is blocking access
- E. The default gateway is incorrect
- F. The local system log file is full

Answer: DE

Explanation:

The most likely reasons why the technician is unable to access a server's package repository internally or externally are that the external firewall is blocking access and that the default gateway is incorrect. A package repository is a source of software packages that can be installed or updated on a server using a package manager tool. A package repository can be accessed over a network using a URL or an IP address. However, if there are any network issues or misconfigurations, the access to the package repository can be blocked or failed. An external firewall is a device or software that controls the incoming and outgoing network traffic based on predefined rules or policies. An external firewall can block access to a package repository if it does not allow traffic on certain ports or protocols that are used by the package manager tool. A default gateway is a device or address that routes network traffic from one network to another network. A default gateway can be incorrect if it does not match the actual device or address that connects the server's network to other networks, such as the internet. An incorrect default gateway can prevent the server from reaching the package repository over other networks.

NEW QUESTION 164

A technician has received tickets responding a server is responding slowly during business hours. Which of the following should the technician implement so the team will be informed of this behavior in real time?

- A. Log rotation
- B. Alerts
- C. Reports
- D. Log stopping

Answer: B

Explanation:

Alerts are notifications that inform the technician or the team of any issues or events that occur on a server or a network. Alerts can be configured to trigger based on certain thresholds, such as CPU usage, disk space, memory utilization, or response time. Alerts can help the technician monitor and troubleshoot the server performance in real time. Verified References: [Alerts], [Server performance]

NEW QUESTION 165

Due to a recent application migration, a company's current storage solution does not meet the necessary requirements for hosting data without impacting performance when the data is accessed in real time by multiple users. Which of the following is the BEST solution for this Issue?

- A. Install local external hard drives for affected users.
- B. Add extra memory to the server where data is stored.
- C. Compress the data to increase available space.
- D. Deploy a new Fibre Channel SAN solution.

Answer: D

Explanation:

A Fibre Channel SAN solution is a type of storage area network (SAN) that uses high-speed optical fiber cables to connect servers and storage devices. A SAN allows for hosting data without impacting performance when the data is accessed in real time by multiple users, as it provides fast data transfer rates, low latency, high availability, and scalability¹². A local external hard drive (A) would not be suitable for multiple users, as it would limit the accessibility and security of the data. Adding extra memory to the server (B) would not solve the problem of data access performance, as it would not increase the bandwidth or reduce the congestion of the network. Compressing the data © would not improve the performance either, as it would add extra overhead and complexity to the data processing and retrieval. References: 1 <https://www.techradar.com/best/best-cloud-storage> 2 <https://solutionsreview.com/data-storage/the-best-enterprise-data-storage-solutions/>

NEW QUESTION 168

Which of the following would a systems administrator implement to ensure all web traffic is secure?

- A. SSH
- B. SSL
- C. SMTP
- D. PGP

Answer: B

Explanation:

Secure Sockets Layer (SSL): SSL and its successor Transport Layer Security (TLS) enable client and server computers to establish a secure connection session and manage encryption and decryption activities. Reference: <https://paginas.fe.up.pt/~als/mis10e/ch8/chpt8-4bullettext.htm>

NEW QUESTION 170

A hardware technician is installing 19 1U servers in a 42U rack. The following unit sizes should be allocated per server?

- A. 1U
- B. 2U
- C. 3U
- D. 4U

Answer: A

Explanation:

1U stands for one unit and it is a standard unit of measurement for rack-mounted servers. It is equal to 1.75 inches (4.45 cm) in height. A 42U rack can accommodate 42 1U servers or a combination of servers with different unit sizes. Therefore, the unit size per server should be 1U if there are 19 1U servers in a 42U rack. References: <https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-objectives> (Objective 1.2)

NEW QUESTION 174

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

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

Answer: B

Explanation:

To ensure the partition is available to the OS, the technician must verify that the volume is formatted as GPT. GPT (GUID Partition Table) is a partitioning scheme that defines how data is organized on a hard disk drive (HDD) or a solid state drive (SSD). GPT uses globally unique identifiers (GUIDs) to identify partitions and supports up to 128 primary partitions per disk. GPT also supports disks larger than 2 TB and has a backup copy of the partition table at the end of the disk for data recovery. GPT is required for installing Windows on UEFI-based PCs, which offer faster boot time and better security than legacy BIOS-based PCs.

NEW QUESTION 177

Which of the following is an architectural reinforcement that is used to attempt to conceal the exterior of an organization?

- A. Fencing
- B. Bollards
- C. Camouflage
- D. Reflective glass

Answer: C

Explanation:

Camouflage is an architectural reinforcement that is used to attempt to conceal the exterior of an organization. Camouflage is a technique of blending in with the surroundings or disguising the appearance of a building or facility to make it less noticeable or identifiable. Camouflage can reduce the visibility and attractiveness of a target for potential attackers or intruders. References: CompTIA Server+ Certification Exam Objectives, Domain 5.0: Security, Objective 5.1: Given a scenario, apply physical security methods to a server.

NEW QUESTION 179

A server administrator is trying to determine the cause of a slowdown on a database server. Upon investigation, the administrator determines the issue is in the storage subsystem. Which of the following will most likely resolve this issue?

- A. Increasing IOPS by implementing flash storage
- B. Implementing deduplication on the storage
- C. Extending capacity by installing a 4TB SATA disk
- D. Reformatting the disk as FAT32

Answer: A

Explanation:

Increasing IOPS (input/output operations per second) by implementing flash storage is the most likely solution to resolve a slowdown issue in the storage subsystem of a database server. Flash storage uses solid-state drives (SSDs) that have faster read/write speeds and lower latency than traditional hard disk drives (HDDs). This can improve the performance of database queries and transactions. Implementing deduplication, extending capacity, or reformatting the disk as FAT32 are not likely to resolve the issue, as they do not affect the IOPS of the storage subsystem. References: [CompTIA Server+ Certification Exam Objectives], Domain 3.0: Storage, Objective 3.5: Summarize hardware and features of various storage technologies.

NEW QUESTION 184

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

- A. 21
- B. 22
- C. 23
- D. 53
- E. 443

F. 636

Answer: AC

Explanation:

The administrator should close ports 21 and 23 to secure the server properly. Port 21 is used for FTP (File Transfer Protocol), which is an unsecure protocol that allows file transfer between a client and a server over a network connection. FTP does not encrypt the data or the credentials that are transmitted, making them vulnerable to interception or modification by attackers. Port 23 is used for Telnet, which is an unsecure protocol that allows remote login and command execution over a network connection using a CLI. Telnet does not encrypt the data or the credentials that are transmitted, making them vulnerable to interception or modification by attackers.

Reference:

<https://www.csoonline.com/article/3191531/securing-risky-network-ports.html>

NEW QUESTION 186

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 191

Due to a disaster incident on a primary site, corporate users are redirected to cloud services where they will be required to be authenticated just once in order to use all cloud services.

Which of the following types of authentications is described in this scenario?

- A. MFA
- B. NTLM
- C. Kerberos
- D. SSO

Answer: D

NEW QUESTION 195

A technician is decommissioning a server from a production environment. The technician removes the server from the rack but then decides to repurpose the system as a lab server instead of decommissioning it. Which of the following is the most appropriate NEXT step to recycle and reuse the system drives?

- A. Reinstall the OS.
- B. Wipe the drives.
- C. Degauss the drives.
- D. Update the IP schema.

Answer: B

Explanation:

Wiping the drives is the most appropriate step to recycle and reuse the system drives. Wiping the drives means erasing all the data on the drives and overwriting them with random or meaningless data. This can help prevent data leakage, comply with regulations, and prepare the drives for a new installation or configuration. Wiping the drives is different from deleting or formatting the drives, which only remove the references to the data but not the data itself. References:

<https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-objectives> (Objective 1.3)

NEW QUESTION 196

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 201

A technician recently upgraded several pieces of firmware on a server. Ever since the technician rebooted the server, it no longer communicates with the network. Which of the following should the technician do FIRST to return the server to service as soon as possible?

- A. Replace the NIC
- B. Make sure the NIC is on the HCL
- C. Reseat the NIC
- D. Downgrade the NIC firmware

Answer: D

Explanation:

The first thing that the technician should do to return the server to service as soon as possible is downgrade the NIC firmware. Firmware is a type of software that controls the basic functions of hardware devices, such as network interface cards (NICs). Firmware updates can provide bug fixes, performance improvements, or new features for hardware devices. However, firmware updates can also cause compatibility issues, configuration errors, or functionality failures if they are not installed properly or if they are not compatible with the device model or driver version. Downgrading the firmware means reverting to an older version of firmware that was previously working fine on the device. Downgrading the firmware can help resolve any problems caused by a faulty firmware update and restore normal operation of the device.

NEW QUESTION 206

Which of the following licensing models allows the greatest number of concurrent Windows VMS to run on a host for the lowest cost?

- A. per user
- B. per core
- C. Per instance
- D. Per concurrent user

Answer: A

Explanation:

The answer to this question may depend on several factors, such as the number and type of Windows VMs, the number and type of host machines, the number and type of users, and the specific licensing terms and conditions of each licensing model. However, based on the information available from the web search results, one possible answer is per user. Per user licensing model is a licensing model that allows a user to access Windows VMs from any device, regardless of the number of devices or VMs. Per user licensing model is available for Windows 10 Enterprise E3/E5, Windows VDA E3/E5, and Microsoft 365 F3/E3/E5. Per user licensing model may offer the greatest number of concurrent Windows VMs to run on a host for the lowest cost if the following conditions are met:

? The user needs to access multiple Windows VMs from different devices, such as desktops, laptops, tablets, or smartphones.

? The user needs to access Windows VMs that run different versions or editions of Windows, such as Windows 10 Enterprise, Windows 10 Pro, or Windows 7 Enterprise.

? The user needs to access Windows VMs that run on different types of host machines, such as physical servers, virtual servers, or cloud servers.

? The user does not need to access Windows VMs that run on dedicated hardware or have specific performance or security requirements.

According to the web search results¹, per user licensing model costs \$84 per user per year for Windows 10 Enterprise E3, \$168 per user per year for Windows 10 Enterprise E5,

\$100.80 per user per year for Windows VDA E3, and \$196.80 per user per year for Windows VDA E5. These prices are based on the Open License Program and may vary depending on the volume and agreement level²

Per core licensing model is a licensing model that requires a license for each core of the processor on the host machine that runs Windows VMs. Per core licensing model is available for Windows Server 2022 Datacenter and Standard editions. Per core licensing model may offer a lower cost than per user licensing model if the following conditions are met:

? The host machine has a low number of cores or a high core density.

? The host machine runs a high number of Windows VMs with low resource consumption.

? The host machine runs only Windows Server VMs with the same edition as the host machine.

According to the web search results², per core licensing model costs \$6,155 for 16 core licenses for Windows Server 2022 Datacenter edition and \$1,069 for 16 core licenses for Windows Server 2022 Standard edition. These prices are suggested retail prices and may vary depending on the reseller²

Per instance licensing model is a licensing model that requires a license for each instance of Windows that runs on a host machine or a VM. Per instance licensing model is available for Windows Server 2022 Essentials edition and some older versions of Windows Server. Per instance licensing model may offer a lower cost than per user or per core licensing model if the following conditions are met:

? The host machine runs only one instance of Windows Server with low resource consumption.

? The host machine does not need to run any other VMs or applications.

? The host machine does not need any advanced features or functions that are available in Datacenter or Standard editions.

According to the web search results², per instance licensing model costs \$501 for one server license for Windows Server 2022 Essentials edition. This price is suggested retail price and may vary depending on the reseller²

Per concurrent user licensing model is a licensing model that allows a certain number of users to access Windows VMs at the same time, regardless of the number of devices or VMs. Per concurrent user licensing model is not available for any current version of Windows or Windows Server. Per concurrent user licensing model was available for some older versions of Windows Server Terminal Services or Remote Desktop Services, but it was discontinued due to complexity and compliance issues. Therefore, per concurrent user licensing model cannot be used for running Windows VMs on a host.

NEW QUESTION 207

A server technician installs a new NIC on a server and configures the NIC for IP connectivity. The technician then tests the connection using the ping command. Given the following partial output of the ping and ipconfig commands:

```
ipconfig /all

IPv4 address: 192.168.1.5
Subnet mask: 255.255.255.0
Default gateway: 192.168.1.1

pinging 192.168.1.1 with 32 bytes of data:

Request timed out
Reply from 192.168.1.1: bytes=32 time<1ms TTL=128
Request timed out
Reply from 192.168.1.1: bytes=32 time<1ms TTL=128
```

Which of the following caused the issue?

- A. Duplicate IP address
- B. Incorrect default gateway
- C. DHCP misconfiguration
- D. Incorrect routing table

Answer: A

Explanation:

? The ping command output shows that the NIC has an IP address of 192.168.1.100 and a default gateway of 192.168.1.1. However, when the technician tries to ping the default gateway, the reply comes from another IP address: 192.168.1.101. This means that there is another device on the network that has the same IP address as the default gateway, and it is responding to the ping request instead of the intended destination.

? A duplicate IP address can cause network connectivity problems, such as packet loss, routing errors, or unreachable hosts. To resolve this issue, the technician should either change the IP address of the default gateway or the device that is conflicting with it, or use DHCP to assign IP addresses automatically and avoid conflicts.

? The other options are not correct because they do not explain the ping output. An incorrect default gateway would cause no reply or a destination unreachable message, not a reply from a different IP address. A DHCP misconfiguration would cause an invalid or no IP address on the NIC, not a duplicate IP address on the network. An incorrect routing table would cause routing errors or unreachable destinations, not a reply from a different IP address.

References:

? https://askleo.com/what_is_ping_and_what_does_its_output_tell_me/

? <https://learn.microsoft.com/en-us/windows-server/administration/windows-commands/ping>

NEW QUESTION 212

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

- A. The server has a faulty power supply.
- B. The server has a CMOS battery failure.
- C. The server requires OS updates.
- D. The server has a malfunctioning LED panel.
- E. The servers have NTP configured.
- F. CPU frequency scaling is set too high.

Answer: BE

Explanation:

A CMOS battery failure can cause the server to lose its BIOS settings, including the date and time, which can affect the server's functionality and connectivity. The servers have NTP (Network Time Protocol) configured to synchronize their clocks with a reliable time source, which can prevent time drift and ensure consistent timestamps. If one server has a wrong date and time, it can cause conflicts and errors with the other servers that have NTP configured.

References:

? CompTIA Server+ Certification Exam Objectives1, page 9

? Signs or symptoms of a CMOS battery failure2

? NTP: Network Time Protocol

NEW QUESTION 215

Hackers recently targeted a company with an attack that resulted in a system breach, which compromised the organization's data. Because of the system breach, the administrator had to bypass normal change management procedures. Which of the following change management activities was necessary?

- A. Cancelled change request
- B. Change request postponement
- C. Emergency change request
- D. Privilege change request
- E. User permission change request

Answer: C

Explanation:

An emergency change request is a type of change request that is initiated in response to an urgent situation, such as a system breach, that requires immediate action to restore normal operations or prevent further damage. An emergency change request may bypass some of the normal change management procedures, such as approval, testing, or documentation, in order to expedite the implementation of the change. However, an emergency change request should still follow the basic steps of change management, such as identification, analysis, planning, execution, and evaluation, and should be reviewed and documented after the change is completed.

References: CompTIA Server+ Study Guide, Chapter 11: Change Management, page 443.

NEW QUESTION 220

A company's IDS has identified outbound traffic from one of the web servers coming over port 389 to an outside address. This server only hosts websites. The company's SOC administrator has asked a technician to harden this server. Which of the following would be the BEST way to complete this request?

- A. Disable port 389 on the server
- B. Move traffic from port 389 to port 443
- C. Move traffic from port 389 to port 637
- D. Enable port 389 for web traffic

Answer: A

Explanation:

The best way to complete the request to harden the server is to disable port 389 on the server. Port 389 is the default port used by LDAP (Lightweight Directory Access Protocol), which is a protocol that allows access and modification of directory services over a network. LDAP can be used for authentication, authorization, or information retrieval purposes. However, LDAP does not encrypt its data by default, which can expose sensitive information or credentials to attackers who can intercept or modify the network traffic.

Therefore, port 389 should be disabled on a web server that only hosts websites and does not need LDAP functionality. Alternatively, port 636 can be used instead of port 389 to enable LDAPS (LDAP over SSL/TLS), which encrypts the data using SSL/TLS certificates.

NEW QUESTION 225

A systems administrator is investigating a server with a RAID array that will not boot into the OS. The administrator notices all the hard drives are reporting to be offline. The administrator checks the RAID controller and verifies the configuration is correct. The administrator then replaces one of the drives with a known-good drive, but it appears to be unavailable as well. Next, the administrator takes a drive out of the server and places it in a spare server, and the drive is available and functional. Which of the following is MOST likely causing the issue?

- A. The kernel is corrupt.
- B. Resources are misallocated.
- C. The backplane has failed.
- D. The drives need to be reseated.

Answer: C

Explanation:

The backplane is a circuit board that connects multiple hard drives to a RAID controller and provides power and data transfer between them. If the backplane has failed, it may cause all the hard drives to be offline and prevent the server from booting into the OS. The fact that replacing one of the drives with a known-good drive did not work, and that taking a drive out of the server and placing it in a spare server made it functional, suggests that the problem is not with the drives themselves but with the backplane. A corrupt kernel (A) would not affect the status of the hard drives, as it is a software component of the OS. Resource misallocation (B) would not cause all the hard drives to be offline, as it is a configuration issue that affects how resources are assigned to processes or applications. Reseating the drives (D) would not help, as it would not fix a faulty backplane. References: <https://www.dell.com/support/kbdoc/en-us/000130114/how-to-troubleshoot-a-faulty-backplane>

NEW QUESTION 229

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

```
dr-xr-xr-- /home/Ann
```

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

- A. `chmod777/home/Ann`
- B. `chmod666/home/Ann`
- C. `chmod711/home/Ann`
- D. `chmod754/home/Ann`

Answer: D

Explanation:

The administrator should use the command `chmod 754 /home/Ann` to resolve the issue without granting unnecessary permissions. The `chmod` command is used to change the permissions of files and directories on a Linux server. The permissions are represented by three numbers, each ranging from 0 to 7, that correspond to the read (r), write (w), and execute (x) permissions for the owner, group, and others respectively. The numbers are calculated by adding up the values of each permission: r = 4, w = 2, x = 1. For example, 7 means rwx (4 + 2 + 1), 6 means rw- (4 + 2), 5 means r-x (4 + 1), etc. In this case, Ann's home directory has the permissions `dr-xr-xr--`, which means that only the owner (d) can read (r) and execute (x) the directory, and the group and others can only read (r) and execute (x) but not write (w) to it. This prevents Ann from saving files to her home directory. To fix this issue, the administrator should grant write permission to the owner by using `chmod 754 /home/Ann`, which means that the owner can read (r), write (w), and execute (x) the directory, the group can read (r) and execute (x) but not write (w) to it, and others can only read (r) but not write (w) or execute (x) it. This way, Ann can save files to her home directory without giving unnecessary permissions to others.

Reference:

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

NEW QUESTION 233

An administrator is troubleshooting performance issues on a server that was recently upgraded. The administrator met with users/stakeholders and documented recent changes in an effort to determine whether the server is better or worse since the changes. Which of the following would BEST help answer the server performance question?

- A. Server performance thresholds
- B. A server baseline
- C. A hardware compatibility list
- D. An application service-level agreement

Answer: B

Explanation:

A server baseline is a set of metrics that represents the normal performance and behavior of a server under a specific workload and configuration. A server baseline can help answer the server performance question by comparing the current performance with the previous performance before the upgrade. This can help identify any changes or issues that may have affected the server performance. References: <https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-objectives> (Objective 4.2)

NEW QUESTION 234

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

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

Answer: C

Explanation:

The administrator should identify the runaway process or processes that are causing high load average on the server. Load average is a metric that indicates how many processes are either running on or waiting for the CPU at any given time. A high load average means that there are more processes than available CPU cores, resulting in poor performance and slow response time. A runaway process is a process that consumes excessive CPU resources without terminating or releasing them. A runaway process can be caused by various factors, such as programming errors, infinite loops, memory leaks, etc. To identify a runaway process, the administrator can use tools such as top, ps, or htop to monitor CPU usage and process status. To stop a runaway process, the administrator can use commands such as kill, pkill, or killall to send signals to terminate it.

NEW QUESTION 237

An administrator is able to ping the default gateway and internet sites by name from a file server. The file server is not able to ping the print server by name. The administrator is able to ping the file server from the print server by both IP address and computer name. When initiating an initiating from the file server for the print server, a different IP address is returned, which of the following is MOST Likely the cause?

- A. A firewall blocking the ICMP echo reply.
- B. The DHCP scope option is incorrect
- C. The DNS entries for the print server are incorrect.
- D. The hosts file misconfigured.

Answer: D

Explanation:

The hosts file is a file that maps hostnames to IP addresses on a server or a computer. It can be used to override or supplement the DNS (Domain Name System) resolution for certain hosts or domains. If the hosts file is misconfigured, it may return a different IP address for a hostname than the one registered in the DNS server, causing connectivity issues or errors. Verified References: [Hosts file], [DNS]

NEW QUESTION 239

A server technician arrives at a data center to troubleshoot a physical server that is not responding to remote management software. The technician discovers the servers in the data center are not connected to a KVM switch, and their out-of-band management cards have not been configured. Which of the following should the technician do to access the server for troubleshooting purposes?

- A. Connect the diagnostic card to the PCIe connector.
- B. Connect a console cable to the server NIC.
- C. Connect to the server from a crash cart.
- D. Connect the virtual administration console.

Answer: C

Explanation:

A crash cart is a mobile device that consists of a monitor, a keyboard, a mouse, and a network connection. It can be used to access a physical server that is not responding to remote management software or does not have out-of-band management cards configured. The technician can connect the crash cart to the server using a VGA or HDMI cable and troubleshoot the server locally. Verified References: [Crash cart], [Out-of-band management]

NEW QUESTION 241

A technician has been tasked to install a new CPU. Prior to the installation the server must be configured. Which of the following should the technician update?

- A. The RAID card
- B. The BIOS
- C. The backplane
- D. The HBA

Answer: B

Explanation:

The BIOS (Basic Input/Output System) is a firmware that controls the initialization and booting of a server. It also provides settings for the CPU, such as speed, voltage, and temperature. Updating the BIOS can improve the performance and compatibility of the CPU and other hardware components. Verified References: [BIOS], [CPU]

NEW QUESTION 245

An administrator is configuring the storage for a new database server, which will host databases that are mainly used for archival lookups. Which of the following storage types

will yield the fastest database read performance?

- A. NAS
- B. SSD
- C. 10K rpm SATA
- D. 15K rpm SCSI

Answer: B

Explanation:

The storage type that will yield the fastest database read performance is SSD. SSD (Solid State Drive) is a type of storage device that uses flash memory to store data. SSDs have no moving parts and can access data faster than traditional hard disk drives (HDDs) that use spinning platters and magnetic heads. SSDs are especially suitable for databases that are mainly used for archival lookups, as they can provide faster response times and lower latency for read operations. References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 1, Lesson 1.2, Objective 1.2

NEW QUESTION 247

A server administrator is setting up a new payroll application. Compliance regulations require that all financial systems logs be stored in a central location. Which of the following should the administrator configure to ensure this requirement is met?

- A. Alerting
- B. Retention
- C. Shipping
- D. Rotation

Answer: C

Explanation:

Shipping is a process of sending logs from one system to another system for centralized storage and analysis. Shipping can help ensure compliance with regulations that require financial systems logs to be stored in a central location. Shipping can also help improve security, performance, and scalability of log management. References: <https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-objectives> (Objective 3.4)

NEW QUESTION 252

Which of the following techniques can be configured on a server for network redundancy?

- A. Clustering
- B. Vitalizing
- C. Cloning
- D. Teaming

Answer: D

Explanation:

Teaming is a technique that can be configured on a server for network redundancy. Teaming involves combining two or more network adapters into a single logical unit that acts as one network interface. This way, if one network adapter fails, another one can take over without disrupting network connectivity. Teaming can also improve network performance by load balancing traffic across multiple network adapters. Clustering is a technique that involves grouping two or more servers together to act as one system for high availability and fault tolerance. Virtualizing is a technique that involves creating multiple virtual machines on a single physical server to optimize resource utilization and flexibility. Cloning is a technique that involves creating an exact copy of a server's configuration and data for backup or migration purposes. References: <https://docs.microsoft.com/en-us/windows-server/networking/technologies/nic-teaming/nic-teaming> <https://www.techopedia.com/definition/19588/clustering> <https://www.techopedia.com/definition/4790/virtualization> <https://www.techopedia.com/definition/4776/cloning>

NEW QUESTION 256

Which of the following access control methodologies can be described BEST as allowing a user the least access based on the jobs the user needs to perform?

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

Answer: B

Explanation:

The access control methodology that can be described best as allowing a user the least access based on the jobs the user needs to perform is role-based access control (RBAC). RBAC is an access control method that assigns permissions to users based on their roles or functions within an organization. RBAC provides fine-grained and manageable access control by defining what actions each role can perform and what resources each role can access. RBAC follows the principle of least privilege, which means that users are only granted the minimum level of access required to perform their tasks. RBAC can reduce security risks, simplify administration, and enforce compliance policies.

NEW QUESTION 258

Which of the following licensing models was created by software companies in response to the increasing density of processors?

- A. Per-instance
- B. Per-server
- C. per-user
- D. per-core

Answer: D

Explanation:

The correct answer is D. per-core.

The per-core licensing model was created by software companies in response to the increasing density of processors. This model is used for software that runs on servers with multi-core processors, and the licensing fee is based on the number of cores. This way, the software vendors can charge more for software that runs on servers with more processing power¹

NEW QUESTION 261

A newly hired systems administrator is concerned about fileshare access at the company. The administrator turns on DLP for the fileshare and lets it propagate for a week. Which of the following can the administrator perform now?

- A. Manage the fileshare from an RDP session.
- B. Audit the permissions of the fileshare.
- C. Audit the access to the physical fileshare.
- D. Manage the permissions from the fileshare.

Answer: B

Explanation:

DLP, or Data Loss Prevention, is a type of security measure that aims to prevent unauthorized access, use, or transfer of sensitive data. DLP can be applied to various types of data, such as email, cloud storage, network traffic, or fileshares¹. DLP for fileshares can help monitor and control who can access, modify, or share files on a network share². By turning on DLP for the fileshare and letting it propagate for a week, the administrator can audit the permissions of the fileshare and see if there are any violations

or anomalies in the access patterns. This can help the administrator identify and remediate any potential risks or compliance issues related to the fileshare². The other options are incorrect because they are not directly related to DLP for fileshares. Managing the fileshare from an RDP session or from the fileshare itself are administrative tasks that do not require DLP. Auditing the access to the physical fileshare is a physical security measure that is not affected by DLP.

NEW QUESTION 263

The management team has mandated the encryption of all server administration traffic. Which of the following should MOST likely be implemented?

- A. SSH
- B. VPN
- C. SELinux
- D. FTPS

Answer: A

Explanation:

SSH stands for Secure Shell and it is a network protocol that provides encrypted and authenticated communication between two hosts. SSH can be used to remotely access and administer a server using a command-line interface or a graphical user interface. SSH can ensure the encryption of all server administration traffic, which can prevent eavesdropping, tampering, or spoofing by unauthorized parties. References: <https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-objectives> (Objective 2.4)

NEW QUESTION 264

.....

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](#)