Exam Questions SK0-005

CompTIA Server+ Certification Exam

https://www.2passeasy.com/dumps/SK0-005/

NEW QUESTION 1
An administrator restores several database files without error while participating in a mock disaster recovery exercise. Later, the administrator reports that the restored databases are
corrupt and cannot be used. Which of the following would best describe what caused this issue?

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

Answer: A

Explanation:
Application consistent backup is a method of backing up data that ensures the integrity and consistency of the application state. It involves notifying the application to flush its data from memory to disk and quiescing any write operations before taking a snapshot of the data. If the databases were not backed up to be application consistent, they might contain incomplete or corrupted data that cannot be restored properly. References:
CompTIA Server+ Certification Exam Objectives1, page 12 What is Application Consistent Backup and How to Achieve It2 Application-Consistent Backups3

NEW QUESTION 2
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 requiredto 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 3
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 highlevel 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 distributesthe 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 4
An administrator is rebooting servers manually after a group of updates were deployed through SCCM. The administrator notices several of the servers did not receive the deployed update. Which of the following should the administrator review first?

A. Confirm the server has the current OS updates and security patches installed.
B. Confirm the server OS has a valid Active Directory account.
C. Confirm the server does not have the firewall running.
D. Confirm the server is in the collection scheduled to receive the update.

Answer: D

Explanation:
The first thing the administrator should check is whether the server is in the collection that was scheduled to receive the update through SCCM. A collection is a group of resources, such as computers or users, that can be managed as a single entity by SCCM. If the server is not in the collection, it will not receive the update. The other options are less likely to be the cause of the problem, as they would affect other aspects of the server’s functionality besides receiving updates.
References: CompTIA Server+ Certification Exam Objectives, Domain 5.0: Security, Objective 5.4: Given a scenario, apply patches/updates and validate their installation.

NEW QUESTION 5

A server administrator is installing a new server that uses 40G0 network connectivity. The administrator needs to find the proper cables to connect the server to the switch. Which of the following connectors should the administrator use?

A. SFP+
B. GBIC
C. SFP
D. QSFP+

Answer: D

Explanation:
QSFP+ is a type of connector that should be used to connect a server to a switch that uses 40G network connectivity. QSFP+ (Quad Small Form-factor Pluggable Plus) is a compact, hot-pluggable transceiver module that supports data rates up to 40 Gbps. QSFP+ modules can be used for various network protocols and media types, such as Ethernet, Fibre Channel, InfiniBand, or optical fiber. QSFP+ modules have a 38-pin edge connector and can be inserted into a QSFP+ port on a switch or a server. SFP+ (Small Form-factor Pluggable Plus) is a type of connector that supports data rates up to 10 Gbps, but not 40 Gbps. SFP+ modules have a 20-pin edge connector and can be inserted into an SFP+ port on a switch or a server. GBIC (Gigabit Interface Converter) is an older type of connector that supports data rates up to 1 Gbps, but not 40 Gbps. GBIC modules have an SC duplex connector and can be inserted into a GBIC port on a switch or a server. SFP (Small Form-factor Pluggable) is another older type of connector that supports data rates up to 1 Gbps or 4 Gbps, but not 40 Gbps. SFP modules have an LC duplex connector and can be inserted into an SFP port on a switch or a server. References: https://www.howtogeek.com/190014/virtualization-basics-understanding-techniques-and- fundamentals/ https://www.howtogeek.com/428483/what-is-end-to-end-encryption-and- why-does-it-matter/ https://www.howtogeek.com/443611/how-to-encrypt-your-macs-system-drive-removable-devices-and-individual-files/

NEW QUESTION 6
A new application server has been configured in the cloud to provide access to all clients within the network. On-site users are able to access all resources, but remote users are reporting issues connecting to the new application. The server administrator verifies that all users are configured with the appropriate group memberships. Which of the following is MOST likely causing the issue?

A. Telnet connections are disabled on the server.
B. Role-based access control is misconfigured.
C. There are misconfigured firewall rules.
D. Group policies have not been applied.

Answer: C

Explanation:
This is the most likely cause of the issue because firewall rules can block or allow traffic based on source, destination, port, protocol, or other criteria. If the firewall rules are not configured properly, they can prevent remote users from accessing the cloud application server, while allowing on-site users to access it.References:https://docs.microsoft.com/en-us/azure/virtual-network/security-overview

NEW QUESTION 7
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 limitthe number of individuals who enter the controlled area and to verify their authorization for physical access1. 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 8
A server technician notices a server is very low on disk space. Upon inspecting the disk utilization, the technician discovers server logs are taxing up a large amount of space. There is no central log server. Which of the following would help free up disk space?

A. Log rotation
B. Log shipping
C. Log alerting
D. Log analysis

Answer: B

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

NEW QUESTION 9
An organization purchased six new 4TB drives for a server. An administrator is tasked with creating an efficient RAID given the minimum disk space requirement of 19TBs. Which of the following should the administrator choose to get the most efficient use of space?

A. RAID 1
B. RAID 5
C. RAID 6
D. RAID 10

Answer: B

Explanation:
RAID 5 is a RAID level that uses disk striping with parity. It requires a minimum of three disks and can handle one disk failure. RAID 5 distributes the parity information across all the disks in the array, which improves the read performance and reduces the write penalty. The capacity of a RAID 5 array is (N-1) times the size of the smallest disk, where N is the number of disks in the array. Therefore, for six 4TB disks, the capacity of a RAID 5 array would be (6-1) x 4TB = 20TB, which meets the minimum disk space requirement of 19TB. RAID 5 also has the leastamount of disk space lost to RAID overhead among the options, as it only uses onedisk's worth of space for parity

NEW QUESTION 10
DRAG DROP
A recent power Outage caused email services to go down. A sever administrator also received alerts from the datacenter's UPS.
After some investigation, the server administrator learned that each POU 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

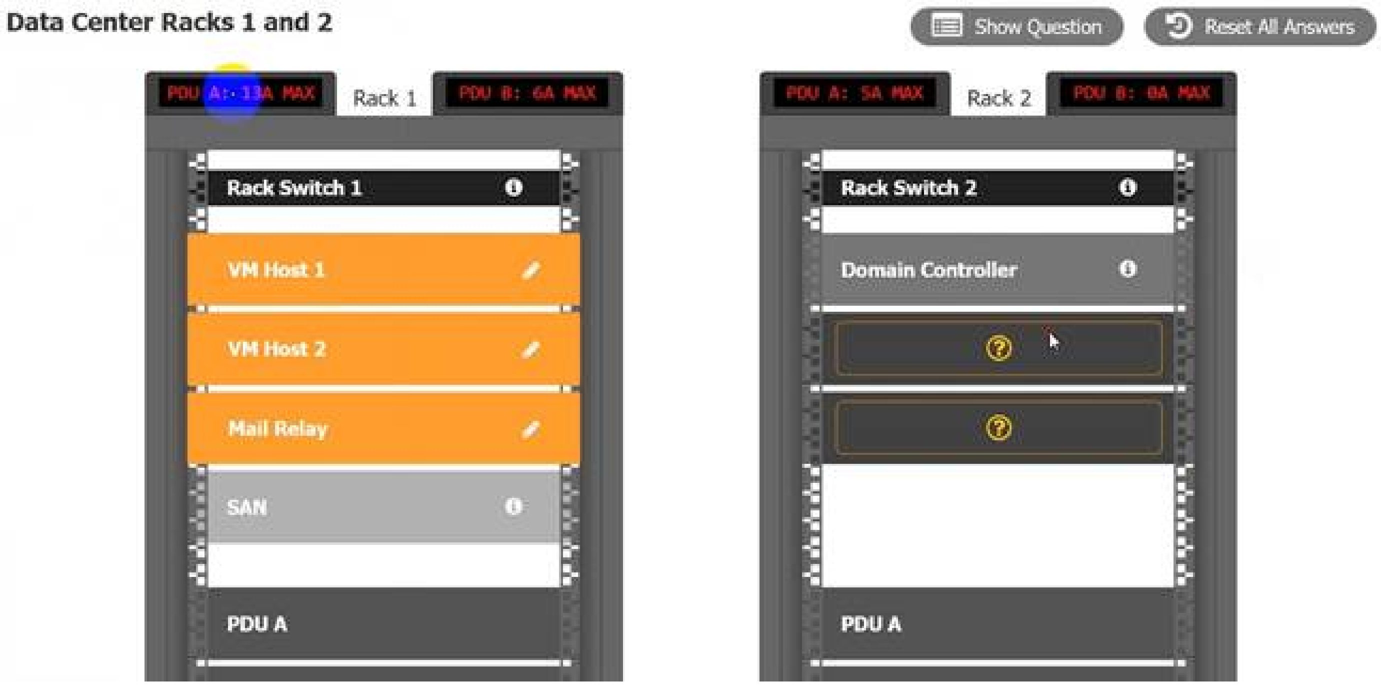

A. Mastered
B. Not Mastered

Answer: A

Explanation:

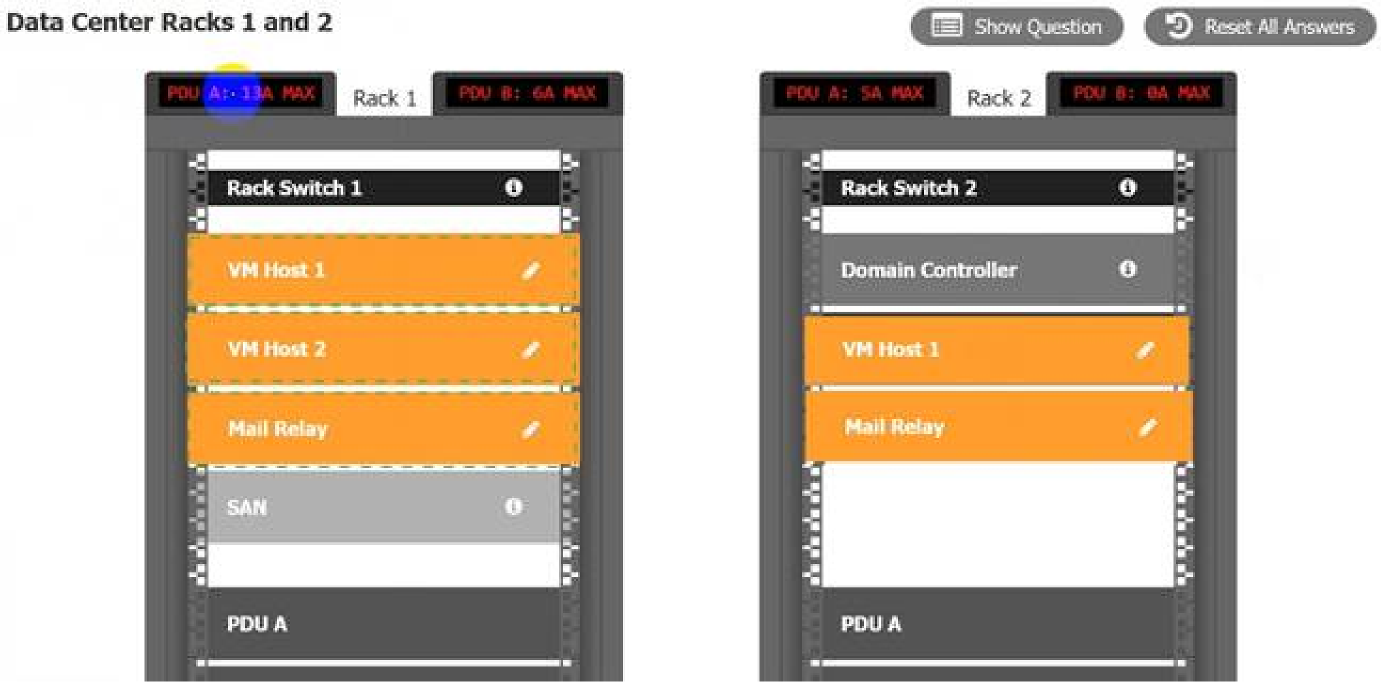

NEW QUESTION 10
A server administrator is installing an OS on a new server. Company policy states no one is to log in directly to the server. Which of the following Installation methods is BEST suited to meet the company policy?

A. GUI
B. Core
C. Virtualized
D. Clone

Answer: B

Explanation:
A core installation is a type of installation method that is best suited to meet the company policy that states no one is to log in directly to the server. A core installation is a minimal installation option that is available when deploying some editions of Windows Server. A core installation includes most but not all server roles and features, but does not include a graphical user interface (GUI). A core installation can only be managed remotely using command-line tools such as PowerShell or Windows Admin Center, or using graphical tools such as Server Manager or Remote Desktop from another computer. This reduces the attack surface, resource consumption, and maintenance requirements of the server. A GUI installation is a type of installation method that includes a graphical user interface (GUI) and allows local or remote management using graphical tools or command- line tools. A virtualized installation is a type of installation method that involves creating and running one or more virtual machines on a physical host using a hypervisor such as Hyper- V or VMware. A clone installation is a type of installation method that involves creating an exact copy of an existing server's configuration and data on another server using tools such as Sysprep or Clonezilla.
References: https://www.howtogeek.com/67469/the- beginners-guide-to-shell-scripting-the-basics/ https://www.howtogeek.com/443611/how-to- encrypt-your-macs-system-drive-removable-devices-and-individual-files/
https://www.howtogeek.com/199068/how-to-upgrade-your-existing-hard-drive-in-under-an- hour/

NEW QUESTION 14
The management team has mandated the use of data-at-rest encryption for all data. Which of the following forms of encryption best achieves this goal?

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

Answer: A

Explanation:
Drive encryption is a form of data-at-rest encryption that encrypts the entire hard drive or solid state drive. This means that all the data on the drive, including the operating system, applications, and files, are protected from unauthorized access. Drive encryption is usually implemented at the hardware or firmware level, and requires a password, PIN, or biometric authentication to unlock the drive. Drive encryption is the most comprehensive and secure way to achieve data-at-rest encryption, as it prevents anyone from accessing the data without the proper credentials, even if they physically remove the drive from the server.
References: CompTIA Server+ Study Guide, Chapter 9: Security, page 367.

NEW QUESTION 15
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 16
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 21
A technician is laying out a filesystem on a new Linux server. Which of the following tools would work BEST to allow the technician to increase a partition's size in the future without reformatting it?

A. LVM
B. DiskPart
C. fdisk

D. Format

Answer: A

Explanation:
LVM (Logical Volume Manager) is a tool that allows the technician to increase a partition’s size in the future without reformatting it on a Linux server. LVM creates logical volumes that can span across multiple physical disks or partitions and can be resized dynamically without losing data. LVM also provides other features such as snapshots, encryption, and RAID. DiskPart, fdisk, and Format are tools that can be used to partition and format disks, but they do not allow increasing a partition’s size without reformatting it. References: https://www.howtogeek.com/howto/40702/how-to-manage-and- use-lvm-logical-volume-management-in-ubuntu/ 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 26
A server administrator needs to harden a server by only allowing secure traffic and DNS inquiries. A port scan reports the following ports are open:

A. 21
B. 22
C. 23
D. 53
E. 443
F. 636

Answer: D

Explanation:
The administrator should only allow secure traffic and DNS inquiries on the server, which means that only ports 22, 53, and 443 should be open. Port 22 is used for SSH (Secure Shell), which is a protocol that allows secure remote login and command execution over a network connection using a command-line interface (CLI). Port 53 is used for DNS (Domain Name System), which is a service that translates domain names into IP addresses and vice versa. Port 443 is used for HTTPS (Hypertext Transfer Protocol Secure), which is a secure version of HTTP that encrypts the data exchanged between a web browser and a web server.
Reference: https://tools.cisco.com/security/center/resources/dns_best_practices

NEW QUESTION 30
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 management activity that is used to address urgent issues that pose a significant risk to the organization, such as a system breach. An emergency change request requires immediate action and approval, and it may bypass some of the normal change management procedures, such as testing, documentation, or stakeholder communication12.
References = 1: Change Management Plans: A Definitive Guide -Indeed(https://www.indeed.com/career-advice/career-development/change-management-activities) 2: The 10 Best Change Management Activities-Connecteam(https://connecteam.com/top-10-change-management-activities/)

NEW QUESTION 34
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 38
An administrator receives an alert stating a S.MAR.T. error has been detected. Which of the following should the administrator run FIRST to determine the issue?

A. A hard drive test
B. A RAM test
C. A power supply swap
D. A firmware update

Answer: A

Explanation:

A S.M.A.R.T. error is an indication of a potential failure of a hard drive.
S.M.A.R.T. stands for Self-Monitoring, Analysis and Reporting Technology and it is a feature that monitors the health and performance of hard drives. A hard drive test can help diagnose the issue and determine if the drive needs to be replaced.References: https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam- objectives (Objective 1.1)

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 45
A Linux server was recently updated. Now, the server stops during the boot process with a blank screen and an£s> prompt. When of the following is the MOST likely cause of this issue?

A. The system is booting to a USB flash drive
B. The UEFI boot was interrupted by a missing Linux boot file
C. The BIOS could not find a bootable hard disk
D. The BIOS firmware needs to be upgraded

Answer: B

Explanation:
The most likely cause of this issue is that the UEFI boot was interrupted by a missing Linux
boot file, such as grub.cfg or vmlinuz, which are essential for loading the Linux kernel and booting the system. The £s> prompt indicates that the system entered into UEFI Shell mode, which is a command-line interface for troubleshooting UEFI boot issues. The administrator can use UEFI Shell commands to locate and restore the missing boot file or change the boot order. Verified References: [UEFI Shell Guide]

NEW QUESTION 50
An administrator has been asked to disable CPU hyperthreading on a server to satisfy a licensing issue. Which of the following best describes how the administrator will likely perform this action?

A. Use a RDP/VNC session.
B. Modify the startup configuration.
C. Use a PowerShell/Bash script.
D. Use the BIOS/UEFI setup.

Answer: D

Explanation:
The BIOS (Basic Input/Output System) or UEFI (Unified Extensible Firmware Interface) setup is a program that allows users to configure the hardware settings of a computer, such as the CPU, memory, disk, and boot options. The BIOS/UEFI setup can be accessed by pressing aspecific key (such as F2, F10, or Delete) during the boot process, before the operating system loads12.
One of the settings that can be changed in the BIOS/UEFI setup is the CPU hyperthreading option. Hyperthreading is a technology that enables a single physical CPU core to execute two threads or tasks simultaneously, improving the performance and efficiency of multi- threaded applications. However, some software licenses may limit the number of CPU cores or threads that can be used, and therefore require disabling hyperthreading on the server34.
To disable hyperthreading on a server, the administrator will likely need to enter the BIOS/UEFI setup and navigate to the processor options menu. There, the administrator will find a setting for Intel ® Hyperthreading Technology or Hyperthreading Function, which can be enabled or disabled. The administrator will need to disable this setting and save the changes. This will turn off hyperthreading on the server and reduce the number of logical CPUs to match the number of physical cores5.

NEW QUESTION 51
An administrator is deploying a new secure web server. The only administration method that is permitted is to connect via RDP. Which of the following ports should be allowed? (Select TWO).

A. 53
B. 80
C. 389
D. 443
E. 45
F. 3389
G. 8080

Answer: DF

Explanation:
Port 443 is the default port for HTTPS, which is the protocol used for secure web communication. HTTPS uses SSL/TLS certificates to encrypt the data between the web server and the browser. Port 443 is commonly used for web servers that need to provide secure services, such as online banking, e-commerce, or email.

By allowing port 443, the administrator can access the web server's interface and manage its settings1.
Port 3389 is the default port for RDP, which is the protocol used for remote desktop connection. RDP allows a user to remotely access and control another computer over a network. Port 3389 is commonly used for remote administration, technical support, or remote work. By allowing port 3389, the administrator can connect to the web server's desktop and perform tasks that require graphical user interface2.

NEW QUESTION 52
Which of the following are measures that should be taken when a data breach occurs? (Select TWO).

A. Restore the data from backup.
B. Disclose the incident.
C. Disable unnecessary ports.
D. Run an antivirus scan.
E. Identify the exploited vulnerability.
F. Move the data to a different location.

Answer: BE

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

NEW QUESTION 53
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 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
An administrator is configuring a server to communicate with a new storage array. To do so, the administrator enters the WWPN of the new array in the server's storage configuration. Which of the following technologies is the new connection using?

A. iSCSI
B. eSATA
C. NFS
D. FcoE

Answer: A

Explanation:
Reference:https://docs.oracle.com/cd/E26996_01/E18549/html/BABHBFHA.html

NEW QUESTION 58
After configuring IP networking on a newly commissioned server, a server administrator installs a straight- through network cable from the patch panel to the switch. The administrator then returns to the server to test network connectivity using the ping command. The partial output of the ping and ipconfig commands are displayed below:

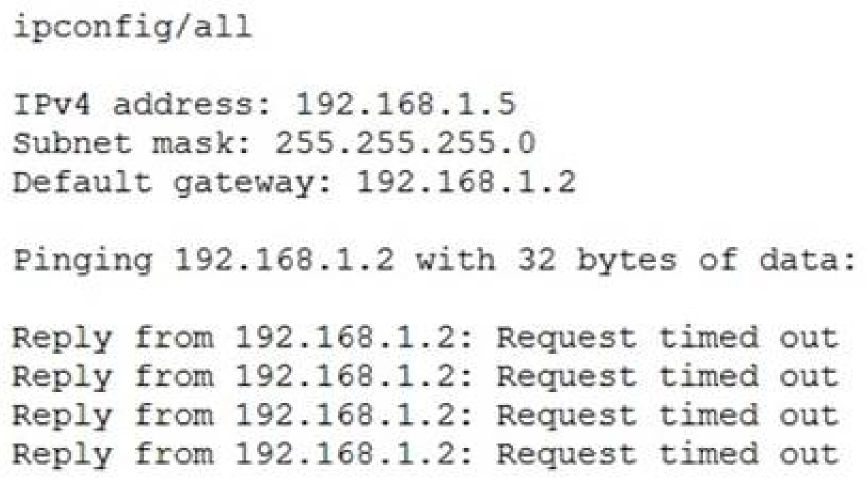

```
ipconfig/all

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

Pinging 192.168.1.2 with 32 bytes of data:

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

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

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

Answer: D

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

NEW QUESTION 60
An administrator is tasked with building an environment consisting of four servers that can each serve the same website. Which of the following concepts is described?

A. Load balancing
B. Direct access
C. Overprovisioning
D. Network teaming

Answer: A

Explanation:
Load balancing is a concept that distributes the workload across multiple servers or other resources to optimize performance, availability, and scalability. Load balancing can be implemented at different layers of the network, such as the application layer, the transport layer, or the network layer. Load balancing can use various algorithms or methods to determine how to distribute the traffic, such as round robin, least connections, or weighted distribution.
References: CompTIA Server+ Study Guide, Chapter 6: Networking, page 241.

NEW QUESTION 64
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 thatstripes 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 66
A security technician generated a public/private key pair on a server. The technician needs to copy the key pair to another server on a different subnet. Which of the following is the most secure method to copy the keys?
? HTTP

A. FTP
B. SCP
C. USB

Answer: C

Explanation:
SCP (Secure Copy Protocol) is a protocol that allows users to securely transfer files between servers using SSH (Secure Shell) encryption. SCP encrypts both the data and the authentication information, preventing unauthorized access, interception, ormodification of the files1. SCP also preserves the file attributes, such as permissions, timestamps, and ownership2.

NEW QUESTION 70
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 73
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/etcdirector
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_GR UB

NEW QUESTION 76
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 withname 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 78
Users at a company work with highly sensitive data. The security department implemented an administrative and technical control to enforce least-privilege access assigned to files. However, the security department has discovered unauthorized data exfiltration. Which of the following is the BEST way to protect the data from leaking?

A. Utilize privacy screens.
B. Implement disk quotas.

C. Install a DLP solution.
D. Enforce the lock-screen feature.

Answer: C

Explanation:
Components of a Data Loss Solution Reference:https://www.imperva.com/learn/data-security/data-loss-prevention-dlp/
The best way to protect the data from leaking is to install a DLP solution. A DLP (Data Loss Prevention) solution is a software that helps businesses prevent confidential data from being leaked or stolen by unauthorized parties. A DLP solution can identify, monitor, and protect data as it moves across networks and devices, such as endpoints, email, web, cloud applications, or removable media. A DLP solution can also enforce security policies based on content and context for data in use, in motion, and at rest. A DLP solution can detect and prevent data breaches by using various techniques, such as content inspection, contextual analysis, encryption, blocking, alerting, warning, quarantining, or other remediation actions.

NEW QUESTION 79
A server administrator receives the following output when trying to ping a local host:

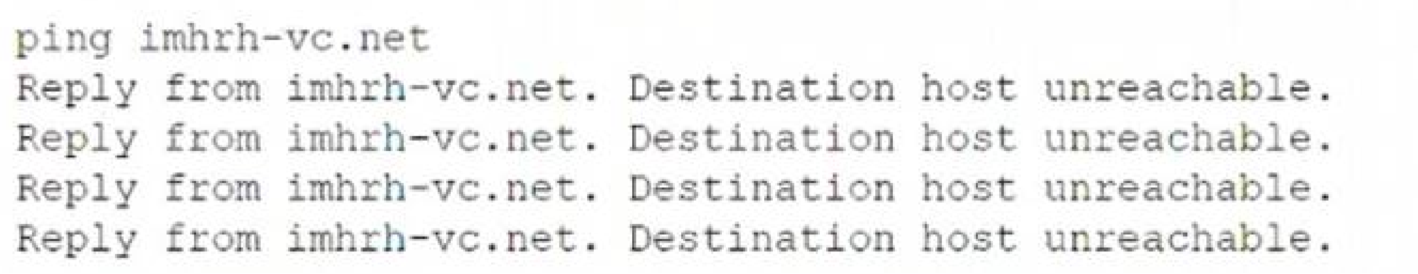

```
ping imhrh-vc.net
Reply from imhrh-vc.net. Destination host unreachable.
Reply from imhrh-vc.net. Destination host unreachable.
Reply from imhrh-vc.net. Destination host unreachable.
Reply from imhrh-vc.net. Destination host unreachable.
```

Which of the following is MOST likely the issue?

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

Answer: A

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

NEW QUESTION 80
An administrator has been asked to verify that all traffic egressing from a company is secured. The administrator confirms all the information that is sent over the network is encrypted. Which of the following describes the type of traffic being encrypted?

A. Network encapsulation
B. Off-site data
C. Secure FTP
D. Data in transit

Answer: D

Explanation:
Data in transit is data that is being transferred over a network, such as the internet. It can be encrypted to protect it from unauthorized access or tampering. Verified References: [Data in transit], [Encryption]

NEW QUESTION 83
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 86
A technician has been asked to check on a SAN. Upon arrival, the technician notices the red LED indicator shows a disk has failed. Which of the following should the technician do NEXT, given the disk is hot swappable?

A. Stop sharing the volume
B. Replace the disk
C. Shut down the SAN
D. Stop all connections to the volume

Answer: B

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

NEW QUESTION 89
A technician recently replaced a NIC that was not functioning. Since then, no device driver is found when starting the server, and the network card is not functioning. Which of the following should the technician check first?

A. The boot log
B. The BIOS
C. The HCL
D. The event log

Answer: C

Explanation:
The technician should check the hardware compatibility list (HCL) first to see if the new NIC is supported by the server's operating system. The HCL is a list of hardware devices that have been tested and verified to work with a specific operating system. If the NIC is not on the HCL, it means that there is no device driver available or compatible for it, and the NIC will not function properly.
References: CompTIAServer+ SK0-005 Certification Study Guide, Chapter 5, Lesson 5.2, Objective 5.2

NEW QUESTION 92
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 96
In which of the following media rotation schemes are daily, weekly, and monthly backup media utilized in a first-in, first-out method?

A. Waterfall
B. Synthetic full
C. Tower of Hanoi
D. Grandfather-father-son

Answer: D

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

NEW QUESTION 100
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 103
A server administrator must respond to tickets within a certain amount of time. The server administrator needs to adhere to the:

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

Answer: D

Explanation:
The server administrator needs to adhere to the Service Level Agreement (SLA) when responding to tickets within a certain amount of time. An SLA is a contract between a service provider and a customer that defines the quality, availability, and responsibilities of the service. An SLA may specify the response time for tickets, as well as other metrics such as uptime, performance, security, and backup frequency.Reference: https://www.ibm.com/cloud/learn/service-level-agreements

NEW QUESTION 108
Which of the following licenses would MOST likely include vendor assistance?

A. Open-source
B. Version compatibility
C. Subscription
D. Maintenance and support

Answer: D

Explanation:
Maintenance and support is a type of license that would most likely include vendor assistance. Maintenance and support is a contract that defines the level and scope of service and assistance that a vendor provides to a customer for using their software product. Maintenance and support may include technical support, bug fixes, patches, updates, upgrades, documentation, training, and other benefits. Maintenance and support licenses usually have an annual fee based on the number of users or devices covered by the contract. Open-source is a type of license that allows free access to the source code and modification and distribution of the software product, but does not guarantee vendor assistance. Version compatibility is not a type of license, but a feature that ensures software products can work with different versions of operating systems or other software products. Subscription is a type of license that allows access to software products for a limited period of time based on recurring payments, but does not necessarily include vendor assistance.References: https://www.techopedia.com/definition/1440/software-licensinghttps://www.techopedia.com/definition/1032/business-impact-analysis-bia

NEW QUESTION 109
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 114
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 115
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 120
A server administrator wants to ensure a storage array can survive the failure of two drives without the loss of data. Which of the following RAID levels should the administrator choose?

A. 1
B. 5
C. 6

Answer: D

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

NEW QUESTION 125
Which of the following physical security concepts would most likely be used to limit personnel access to a restricted area within a data center?

A. An access control vestibule
B. Video surveillance
C. Bollards
D. Data center camouflage

Answer: A

Explanation:
An access control vestibule is a physical security concept that limits personnel access to a restricted area within a data center. It is a small room or hallway that has two doors: one that leads to the outside and one that leads to the restricted area. The doors are controlled by an electronic lock that requires authentication, such as a card reader, biometric scanner, or keypad. Only authorized personnel can enter the vestibule and access the restricted area. 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 126
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 127
A server administrator needs to create a new folder on a file server that only specific users can access. Which of the following BEST describes how the server administrator can accomplish this task?

A. Create a group that includes all users and assign it to an ACL.
B. Assign individual permissions on the folder to each user.C Create a group that includes all users and assign the proper permissions.
C. Assign ownership on the folder for each user.

Answer: C

Explanation:
The top portion of the dialog box lists the users and/or groups that have access to the file or folder.
Reference:https://www.uwec.edu/kb/article/drives-establishing-windows-file-and-folder- level-permissions/

NEW QUESTION 131
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 thedatabasefrom 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 I 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 135
A company is building a new datacenter next to a busy parking lot. Which of the following is the BEST strategy to ensure wayward vehicle traffic does not interfere with datacenter operations?

A. Install security cameras
B. Utilize security guards
C. Install bollards
D. Install a mantrap

Answer: C

Explanation:
The best strategy to ensure wayward vehicle traffic does not interfere with datacenter operations is to install bollards. Bollards are sturdy posts that are installed around a perimeter to prevent vehicles from entering or crashing into a protected area. Bollards can provide physical security and deterrence for datacenters that are located near busy roads or parking lots. Bollards can also prevent accidental damage or injury caused by vehicles that lose control or have faulty brakes.

NEW QUESTION 138
Which of the following BEST describes overprovisioning in a virtual server environment?

A. Committing more virtual resources to virtual machines than there are physical resources present
B. Installing more physical hardware than is necessary to run the virtual environment toallow for future expansion
C. Allowing a virtual machine to utilize more resources than are allocated to it based on the server load
D. Ensuring there are enough physical resources to sustain the complete virtual environment in the event of a host failure

Answer: A

Explanation:
This is the best definition of overprovisioning in a virtual server environment because it means allocating more CPU, memory, disk, or network resources to the virtual machines than what is actually available on the physical host. This can lead to performance issues and resource contention.
References:https://www.hpe.com/us/en/insights/articles/10-virtualization-mistakes-everyone-makes-1808.html

NEW QUESTION 140
The Chief Information Officer (CIO) of a datacenter is concerned that transmissions from the building can be detected from the outside. Which of the following would resolve this concern? (Choose two.)

A. RFID
B. Proximity readers
C. Signal blocking

D. Camouflage
E. Reflective glass
F. Bollards

Answer: CE

Explanation:
The best solutions to resolve the concern of transmissions from the building being detected from outside are signal blocking and reflective glass. Signal blocking is a method of preventing or interfering with electromagnetic signals from escaping or entering a certain area. Signal blocking can be achieved by using various materials or devices that create physical barriers or generate noise or jamming signals. Signal blocking can protect data transmissions from being intercepted or eavesdropped by unauthorized parties. Reflective glass is a type of glass that has a coating or film that reflects light and heat. Reflective glass can reduce glare and solar radiation, as well as prevent visual observation from outside. Reflective glass can enhance privacy and security for datacenter operations.

NEW QUESTION 142
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 theolder 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 143
An administrator notices nigh traffic on a certain subnetand wouldlike to identify the source of the traffic. Which of the following tools should the administrator utilize?

A. Anti-malware
B. Nbtstat
C. Port scanner
D. Sniffer

Answer: D

Explanation:
A sniffer is a tool that captures and analyzes network traffic on a subnet or a network interface. It can help identify the source, destination, protocol, and content of the traffic and detect any anomalies or issues on the network. Verified References: [Sniffer], [Network traffic]

NEW QUESTION 145
Which of the following BEST measures now much downtime an organization can tolerate Curing 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- biahttps://www.techopedia.com/definition/8239/mean-time-to-repair-mttr

NEW QUESTION 146
A technician is connecting a server’s secondary NIC to a separate network. The technician connects the cable to the switch but then does not see any link lights on the NIC. The technician confirms there is nothing wrong on the network or with the physical connection. Which of the following should the technician perform NEXT?

A. Restart the server
B. Configure the network on the server
C. Enable the port on the server
D. Check the DHCP configuration

Answer: C

Explanation:
The next thing that the technician should perform is to enable the port on the server. A port is a logical endpoint that identifies a specific service or application on a network device. A port can be enabled or disabled depending on whether the service or application is running or not. If a port is disabled on a server, it means that the server cannot send or receive any network traffic on that port, which can prevent communication with other devices or services that use that port. In this case, if port 389 is disabled on the server, it means that the server cannot use LDAP to access or modify directory services over a network. To resolve this issue, the technician should enable port 389 on the server using commands such as netsh or iptables.

NEW QUESTION 149
A security analyst completed a port scan of the corporate production-server network. Results of the scan were then provided to a systems administrator for immediate action. The following table represents the requested changes:

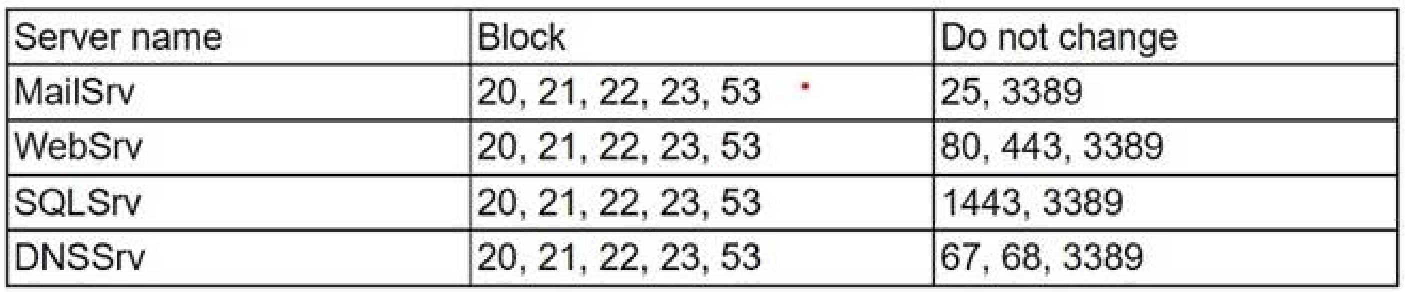

Server name	Block	Do not change
MailSrv	20, 21, 22, 23, 53	25, 3389
WebSrv	20, 21, 22, 23, 53	80, 443, 3389
SQLSrv	20, 21, 22, 23, 53	1443, 3389
DNSSrv	20, 21, 22, 23, 53	67, 68, 3389

The systems administrator created local firewall rules to block the ports indicated above. Immediately, the service desk began receiving calls about the internet being down. The systems administrator then reversed the changes, and the internet became available again. Which of the following ports on DNSSrv must remain open when the firewall rules are reapplied?

A. 20
B. 21
C. 22
D. 23
E. 53

Answer: E

Explanation:
Port 53 is the standard port for DNS (Domain Name System) queries and responses. DNS is a service that translates domain names (such as www.example.com) into IP addresses (such as 192.0.2.1) and vice versa. DNS is essential for internet connectivity, as it allows users and applications to access websites and other online resources by using human- readable names instead ofnumerical addresses1.
The DNSSrv server is a DNS server that provides name resolution for the corporate network. If port 53 is blocked on this server, it will not be able to communicate with other DNS servers or clients, and the name resolution will fail. This will prevent users from accessing any websites or online services that rely on domain names, such as web browsers, email clients, or cloud applications. Therefore, port 53 must remain open on DNSSrv to allow DNS traffic to flow.

NEW QUESTION 151
Several new components have been added to a mission-critical server, and corporate policy states all new components must meet server hardening requirements. Which of the following should be applied?

A. Definition updates
B. Driver updates
C. OS security updates
D. Application updates

Answer: B

Explanation:
Driver updates should be applied to the new components that have been added to a mission-critical server, as part of the server hardening requirements. Drivers are software programs that enable the communication and functionality of hardware devices, such as network cards, storage controllers, or graphics cards. Updating drivers can improve the performance, compatibility, and stability of the new components with the server operating system and applications. References: CompTIA Server+ Certification Exam Objectives, Domain 2.0: Hardware, Objective 2.2: Given a scenario, install, configure and maintain server components.

NEW QUESTION 154
A site is considered a warm site when it:
? has basic technical facilities connected to it.
? has faulty air conditioning that is awaiting service.
? is almost ready to take over all operations from the primary site.

A. is fully operational and continuously providing services.

Answer: A

Explanation:
A warm site is a backup site that has some of the necessary hardware, software, and network resources to resume operations, but not all of them. A warm site requires some time and effort to become fully operational. A warm site is different from a cold site, which has minimal or no resources, and a hot site, which has all the resources and is ready to take over immediately.
References: CompTIA Server+ Study Guide, Chapter 10: Disaster Recovery, page 403.

NEW QUESTION 155
An administrator discovers a Bash script file has the following permissions set in octal notation;
777
Which of the following is the MOST appropriate command to ensure only the root user can modify and execute the script?

A. chmod go-rw>:
B. chmod u=rwx

C. chmod u+wx
D. chmod g-rwx

Answer: A

Explanation:
chmod is a command-line tool that changes the permissions of files and directories in Linux and Unix systems. chmod go-rwx means to remove read, write, and execute permissions for group and other users from a file or directory. This can ensure only the root user can modify and execute the script, since root user has full access to all files and directories regardless of their permissions. References:https://linux.die.net/man/1/chmod

NEW QUESTION 159
A senior administrator instructs a technician to run the following script on a Linux server: for i in {1..65536}; do echo Si; telnet localhost $i; done
The script mostly returns the following message: Connection refused. However, there are several entries in the console display that look like this:
80
Connected to localhost 443
Connected to localhost
Which of the following actions should the technician perform NEXT?

A. Look for an unauthorized HTTP service on this server
B. Look for a virus infection on this server
C. Look for an unauthorized Telnet service on this server
D. Look for an unauthorized port scanning service on this server.

Answer: A

Explanation:
The script that the technician is running is trying to connect to every port on the localhost (the same machine) using telnet, a network protocol that allows remote access to a command-line interface. The script mostly fails because most ports are closed or not listening for connections. However, the script succeeds on ports 80 and 443, which are the default ports for HTTP and HTTPS protocols, respectively. These protocols are used for web services and web browsers. Therefore, the technician should look for an unauthorized HTTP service on this server, as it may indicate a security breach or a misconfiguration. Looking for a virus infection on this server is also possible, but not the most likely source of the issue. Looking for an unauthorized Telnet service on this server is not relevant, as the script is using telnet as a client, not a server. Looking for an unauthorized port scanning service on this server is not relevant, as the script is scanning ports on the localhost, not on other machines. References:
? https://phoenixnap.com/kb/telnet-windows
? https://www.techopedia.com/definition/23337/http-port-80
? https://www.techopedia.com/definition/23336/https-port-443

NEW QUESTION 160
A company wants to deploy software to all users, Out very few of men will be using the software at any one point in time. Which of the following licensing models would be BEST lot the company?

A. Per site
B. Per concurrent user
C. Per core
D. Per instance

Answer: B

Explanation:
Per concurrent user licensing is a model that allows a fixed number of users to access the software at any one point in time. This model is best for the company that wants to deploy software to all users, but very few of them will be using the software at any one point in time. This way, the company can save money by paying only for the number of simultaneous users, rather than for every user who has access to the software. Per site licensing is a model that allows unlimited users within a specific location to use the software. Per core licensing is a model that charges based on the number of processor cores on the server where the software is installed. Per instance licensing is a model that charges based on the number of copies of the software running on different servers or virtual machines.References: https://www.pcmag.com/encyclopedia/term/concurrent-use-licensehttps://www.techopedia.com/definition/1440/software-licensing

NEW QUESTION 162
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 167
A company stores extremely sensitive data on an alt-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 172
A user can successfully connect to a database server from a home office but is unable to access it from a hotel room. Which of the following authentication methods is most likely
configured?

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

Answer: D

Explanation:
Scope-based authentication is a method of restricting access to resources based on the location, network, or device of the user. It can be used to prevent unauthorized access from outside the organization's network or from untrusted devices. In this case, the user can connect to the database server from the home office, which is likely within the scope of the authentication policy, but not from the hotel room, which is outside the scope.
References:
CompTIA Server+ Certification Exam Objectives1, page 15 CompTIA Server+: Authentication & Authorization2

NEW QUESTION 175
An administrator has been asked to increase the storage capacity of a stand-alone file server but no further expansion slots are available. Whichof the following would be the FASTEST solution to implement with no downtime?

A. Configure a RAID array.
B. Replace the current drives with higher-capacity disks.
C. Implement FCoE for more storage capacity.
D. Connect the server to a SAN

Answer: D

Explanation:
A SAN (Storage Area Network) is a network of storage devices that can provide shared storage capacity to multiple servers. By connecting the server to a SAN, the administrator can increase the storage capacity of the server without adding any internal disks or expansion cards. This solution can be implemented quickly and without any downtime. Verified References: [What is a SAN and how does it differ from NAS?]

NEW QUESTION 176
HOTSPOT
A systems administrator deployed a new web proxy server onto the network. The proxy server has two interfaces: the first is connected to an internal corporate firewall, and the second is connected to an internet-facing firewall. Many users at the company are reporting they are unable to access the Internet since the new proxy was introduced. Analyze the network diagram and the proxy server's host routing table to resolve the Internet
connectivity issues.
INSTRUCTIONS
Perform the following steps:
* 1. Click on the proxy server to display its routing table.
* 2. Modify the appropriate route entries to resolve the Internet connectivity issue.
If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

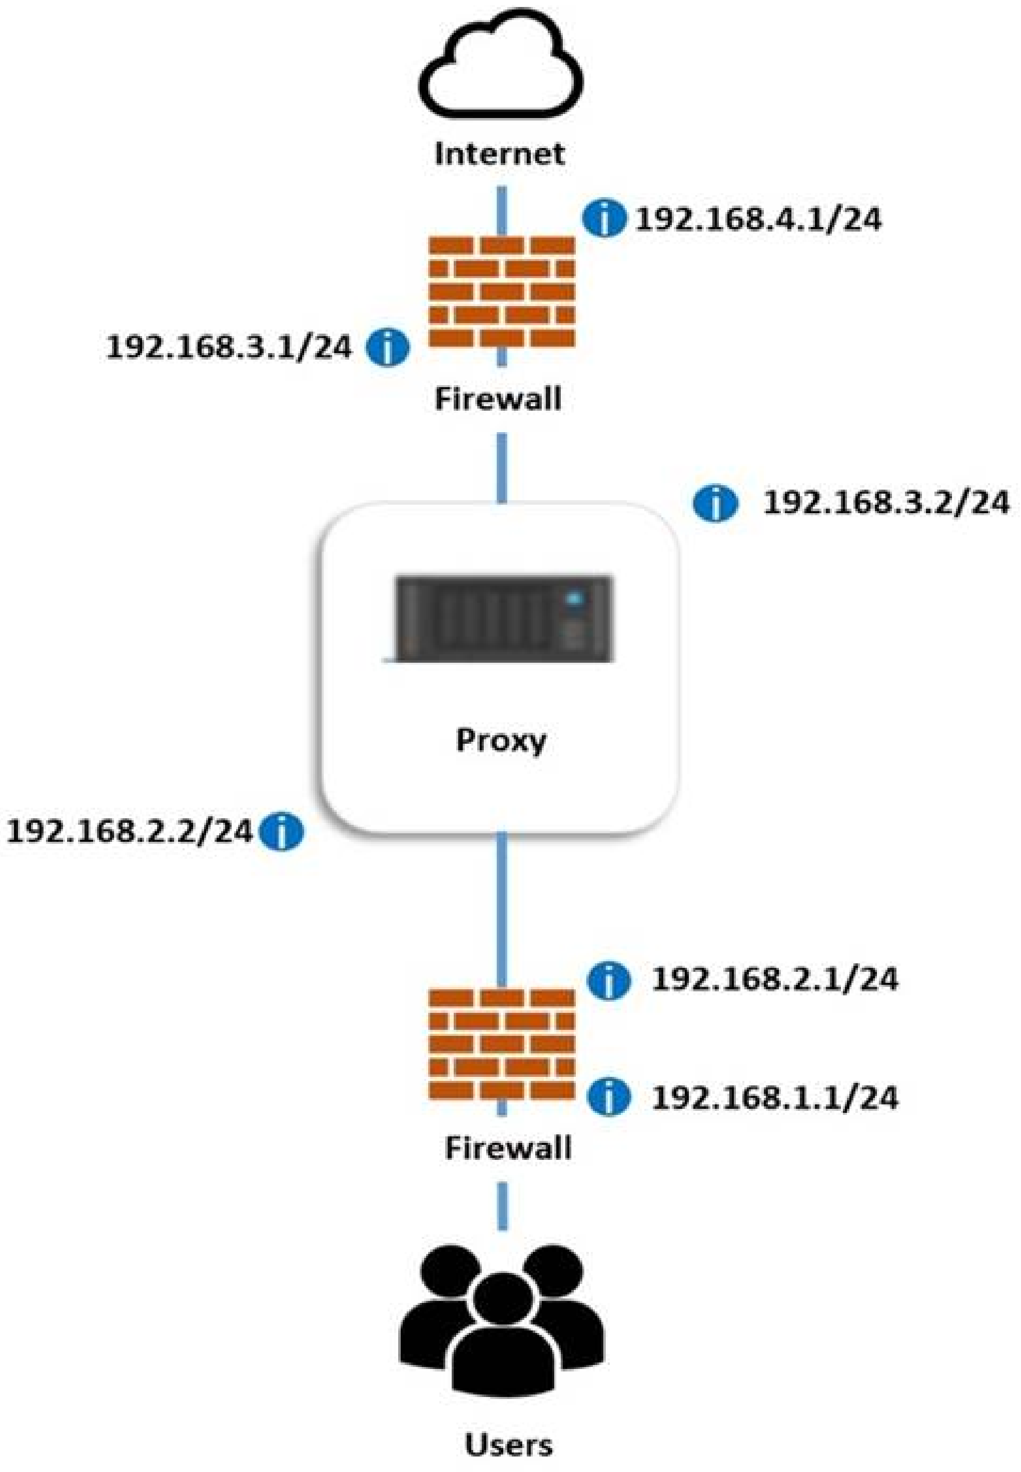

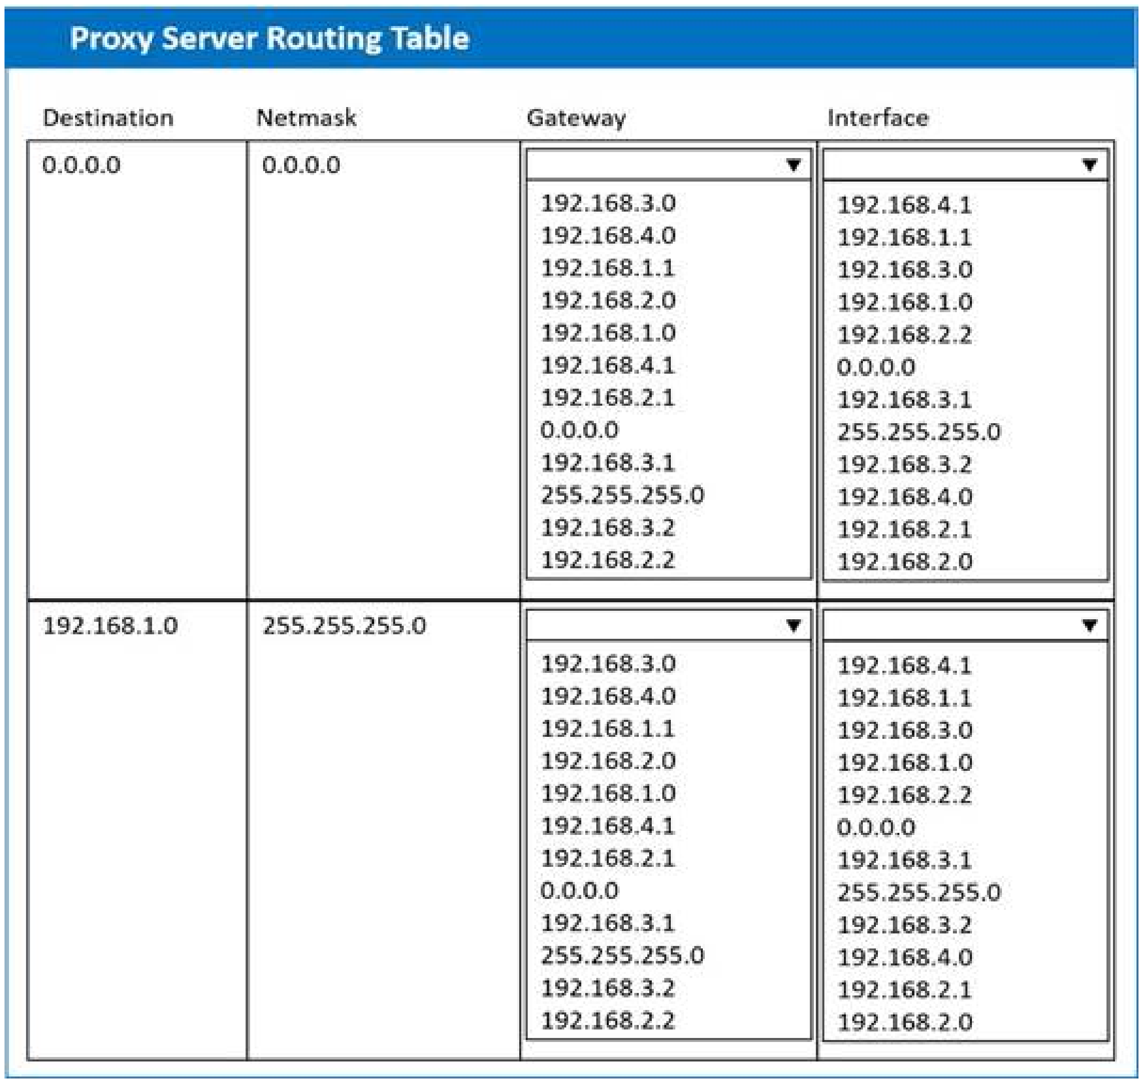

Proxy Server Routing Table

Destination	Netmask	Gateway	Interface
0.0.0.0	0.0.0.0	▼ 192.168.3.0 192.168.4.0 192.168.1.1 192.168.2.0 192.168.1.0 192.168.4.1 192.168.2.1 0.0.0.0 192.168.3.1 255.255.255.0 192.168.3.2 192.168.2.2	▼ 192.168.4.1 192.168.1.1 192.168.3.0 192.168.1.0 192.168.2.2 0.0.0.0 192.168.3.1 255.255.255.0 192.168.3.2 192.168.4.0 192.168.2.1 192.168.2.0
192.168.1.0	255.255.255.0	▼ 192.168.3.0 192.168.4.0 192.168.1.1 192.168.2.0 192.168.1.0 192.168.4.1 192.168.2.1 0.0.0.0 192.168.3.1 255.255.255.0 192.168.3.2 192.168.2.2	▼ 192.168.4.1 192.168.1.1 192.168.3.0 192.168.1.0 192.168.2.2 0.0.0.0 192.168.3.1 255.255.255.0 192.168.3.2 192.168.4.0 192.168.2.1 192.168.2.0

A. Mastered
B. Not Mastered

Answer: A

Explanation:

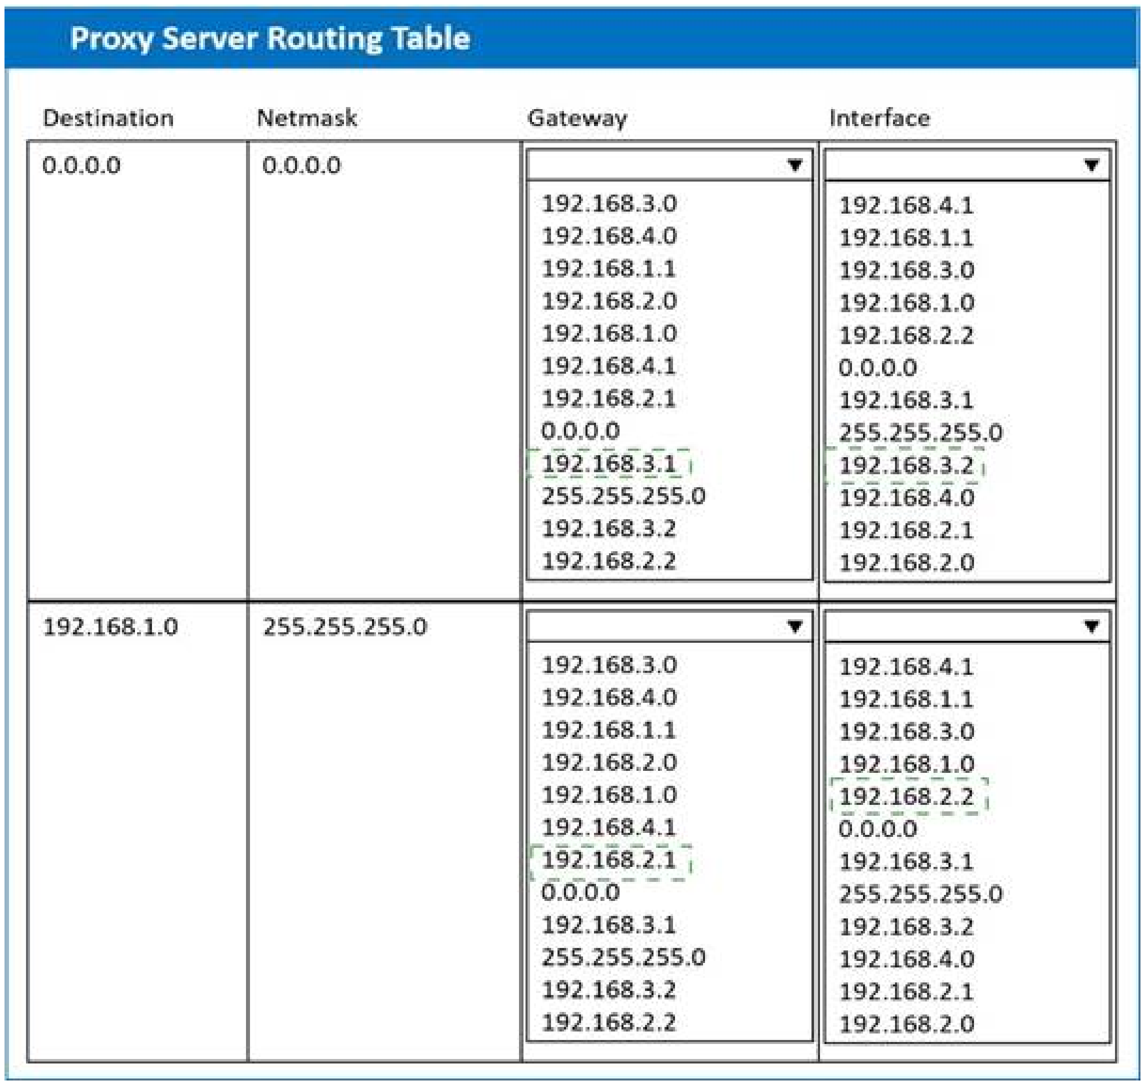

Proxy Server Routing Table

Destination	Netmask	Gateway	Interface
0.0.0.0	0.0.0.0	▼ 192.168.3.0 192.168.4.0 192.168.1.1 192.168.2.0 192.168.1.0 192.168.4.1 192.168.2.1 0.0.0.0 [✓ 192.168.3.1] 255.255.255.0 192.168.3.2 192.168.2.2	▼ 192.168.4.1 192.168.1.1 192.168.3.0 192.168.1.0 192.168.2.2 0.0.0.0 192.168.3.1 255.255.255.0 [✓ 192.168.3.2] 192.168.4.0 192.168.2.1 192.168.2.0
192.168.1.0	255.255.255.0	▼ 192.168.3.0 192.168.4.0 192.168.1.1 192.168.2.0 192.168.1.0 192.168.4.1 [✓ 192.168.2.1] 0.0.0.0 192.168.3.1 255.255.255.0 192.168.3.2 192.168.2.2	▼ 192.168.4.1 192.168.1.1 192.168.3.0 192.168.1.0 [✓ 192.168.2.2] 0.0.0.0 192.168.3.1 255.255.255.0 192.168.3.2 192.168.4.0 192.168.2.1 192.168.2.0

NEW QUESTION 179
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) x Smin, where n is the number of disks and Smin is the smallest disk size. In this case, the RAID 6 capacity is (5-2) x 4TB = 12TB. References:
? CompTIA Server+ Certification Exam Objectives1, page 8
? RAID Levels and Types Explained: Advantages and Disadvantages2
? RAID Levels & Fault Tolerance3

NEW QUESTION 180
A server has experienced several component failures. To minimize downtime, the server administrator wants to replace the components while the server is running. Which of the following can MOST likely be swapped out while the server is still running? (Select TWO).

A. The power supply
B. The CPU
C. The hard drive
D. The GPU
E. The cache
F. The RAM

Answer: AC

Explanation:
The power supply and the hard drive are two components that can most likely be swapped out while the server is still running, if they support hot swapping or hot plugging. Hot swapping or hot plugging means that the device can be added or removed without shutting down the system. The operating system automatically recognizes the changes that have been made. This feature is useful for minimizing downtime and improving availability. The CPU, the GPU, the cache, and the RAM are not hot swappable and require the system to be powered off before replacing them. References: https://www.geeksforgeeks.org/what-is-hot-swapping/https://www.howtogeek.com/268249/what-is-hot-swapping-and-what-devices- support-it/

NEW QUESTION 183
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 187
An administrate 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 192
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 195
Which of the following BEST describes a warm site?

A. The site has all infrastructure and live data.
B. The site has all infrastructure and some data
C. The site has partially redundant infrastructure and no network connectivity
D. The site has partial infrastructure and some data.

Answer: D

Explanation:
A warm site is a type of disaster recovery site that has some pre-installed hardware, software, and network connections, but not as much as a hot site. A warm site also has some backup data, but not as current as a hot site. A warm site requires some time and effort to become fully operational in the event of a disaster. A hot site is a disaster recovery site that has all infrastructure and live data, and can take over the primary site’s operations immediately. A cold site is a disaster recovery site that has no infrastructure or data, and requires significant time and resources to set up. References:
? https://www.enterprisestorageforum.com/management/disaster-recovery-site/
? https://www.techopedia.com/definition/3780/warm-site

NEW QUESTION 199
Which of the following encryption methodologies would MOST likely be used to ensure encrypted data cannot be retrieved if a device is stolen?

A. End-to-end encryption
B. Encryption in transit
C. Encryption at rest
D. Public key encryption

Answer: C

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

NEW QUESTION 203
An administrator is investigating several unexpected documents and video files that recently appeared in a network share. The administrator checks the properties of the files and sees the author's name on the documents is not a company employee. The administrator questions the other users, but no one knows anything about the files. The administrator then checks the log files and discovers the FTP protocol was used to copy the files to the server. Which of the following needs to be done to prevent this from happening again?

A. Implement data loss prevention.
B. Configure intrusion detection.
C. Turn on User Account Control.
D. Disable anonymous access.

Answer: D

Explanation:
This is the best solution to prevent unauthorized files from being copied to the server via FTP because anonymous access allows anyone to log in to the FTP server without providing a username or password. Disabling anonymous access will require users to authenticate with valid credentials before accessing the FTP server.References: https://docs.microsoft.com/en-us/iis/configuration/system.applicationhost/sites/site/ftpserver/security/authentication/anony mousauthentication

NEW QUESTION 206
A data center has 4U rack servers that needto be replaced using VMsbutwithout losingany data. Whichofthefollowingmethodswill MOST likelybe used to replace these servers?

A. Unattended scripted OS installation
B. P2V
C. VM cloning

Answer: C

Explanation:
P2V (Physical to Virtual) is a method of converting a physical server into a virtual machine that can run on a hypervisor. This method can be used to replace 4U rack servers with VMs without losing any data, as it preserves the configuration and state of the original server. P2V can also reduce hardware costs, power consumption, and space requirements. Verified References: [What is P2V?]

NEW QUESTION 207
Due to a recent application migration, a company’s current storage solution does not meet the necessary requirements tor 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 extraoverhead 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 209
A server administrator recently installed a kernel update to test functionality Upon reboot, the administrator determined the new kernel was not compatible with certain server hardware and was unable to uninstall the update. Which of the following should the administrator do to mitigate further issues with the newly instated kernel version?

A. Edit the bootloader configuration file and change the first Kernel stanza to reflect the file location for the last known-good kernel files.
B. Perform a complete OS reinstall on the server using the same media that was used during the initialinstall.
C. Edit the bootloader configuration file and move the newest kernel update stanza lo the end of the file.
D. Set a BIOS password to prevent server technicians from making any changes to thesystem.

Answer: A

Explanation:
The bootloader configuration file is used to specify which kernel version and options to use when booting the system. The first kernel stanza in the file is the default one that is loaded automatically. By editing this stanza and changing it to point to the last known-good kernel files, the administrator can boot the system with a working kernel and avoid any compatibility issues with the new kernel update. Verified References: [How To Change The Linux Kernel Version]

NEW QUESTION 210
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 211
An administrator is installing a new server and OS. After installing the OS, the administrator logs in and wants to quickly check the network configuration. Which of the following is the best command to use to
accomplish this task?

A. tracert
B. telnet
C. ipconfig
D. ping

Answer: C

NEW QUESTION 216
A technicianretailed a new4TBharddrive inaWindows server. Which of the following should the technician perform FIRST to provision the newdrive?

A. Configure the drive as a base disk.
B. Configure the drive as a dynamic disk.
C. Partition the drive using MBR.
D. Partition the drive using OPT.

Answer: D

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. MBR (Master Boot Record) is an older partitioning scheme that has limitations on the size and number of partitions (up to 4 primary partitions or 3 primary and 1 extended partition per drive). To provision a new 4 TB drive, the technician should partition it using GPT. Verified References: [GPT], [MBR]

NEW QUESTION 218
A technician is able to copy a Me to a temporary folder on another partition but is unable to copy it to a network share or a USB flash drive. Which of the following is MOST likely preventing the file from being copied to certain locations?

A. An ACL
B. Antivirus
C. DLP
D. A firewall

Answer: C

Explanation:
DLP (Data Loss Prevention) is a security measure that prevents unauthorized copying, transferring, or leaking of sensitive data from a server or a network. It can block or alert the user when they try to copy a file to certain locations, such as a network share or a USB flash drive, based on predefined policies and rules. Verified References: [DLP], [Data loss]

NEW QUESTION 223
Which of the following types of asset management documentation is commonly used as a reference when processing the replacement of a faulty server component?

A. Warranty
B. Purchase order
C. License

D. Baseline document

Answer: A

Explanation:
A warranty is a type of asset management documentation that is commonly used as a reference when processing the replacement of a faulty server component. A warranty is a guarantee from the manufacturer or vendor that covers the repair or replacement of defective parts within a specified period of time. A purchase order, a license, or a baseline document are not directly related to the replacement of a faulty server component. References: [CompTIA Server+ Certification Exam Objectives], Domain 1.0: Server Architecture, Objective 1.4: Explain asset management and documentation processes.

NEW QUESTION 226
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 227
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 228
Which of the following attacks is the most difficult to mitigate with technology?

A. Ransomware
B. Backdoor
C. SQL injection
D. Phishing

Answer: D

Explanation:
Phishing is a type of attack that is the most difficult to mitigate with technology. Phishing is a technique of deceiving users into revealing their personal or confidential information, such as passwords, credit card numbers, or bank accounts, by sending them fraudulent emails or messages that appear to be from legitimate sources. Phishing relies on human factors, such as curiosity, greed, or fear, to trick users into clicking on malicious links or attachments, or entering their credentials on fake websites. Technology solutions, such as antivirus software, firewalls, or spam filters, can help detect and block some phishing attempts, but they cannot prevent users from falling victim to social engineering tactics. References: [CompTIA Server+ Certification Exam Objectives], Domain 5.0: Security, Objective 5.3: Given a scenario, explain methods and techniques to secure data.

NEW QUESTION 229
A company is reviewing options for its current disaster recovery plan and potential changes to it. The security team will not allow customer data to egress to non-company equipment, and the company has requested recovery in the shortest possible time. Which of the following will BEST meet these goals?

A. A warm site
B. A hot site
C. Cloud recovery
D. A cold site

Answer: B

Explanation:
A hot site is a type of disaster recovery site that has all the equipment and data ready to resume operations as soon as possible after a disaster. A hot site is usually located in a different geographic area than the primary site and has redundant power, cooling, network, and security systems. A hot site is best for the company that wants to recover in the shortest possible time and does not want customer data to egress to non- company equipment. A warm site is a type of disaster recovery site that has some equipment and data ready, but requires some configuration and restoration before resuming operations. A cold site is a type

of disaster recovery site that has only basic infrastructure and space available, but requires significant setup and installation before resuming operations. Cloud recovery is a type of disaster recovery service that uses cloud- based resources and platforms to store backups and restore data and applications after a disaster. References: https://www.techopedia.com/definition/11172/hot-site https://www.techopedia.com/definition/11173/warm-site https://www.techopedia.com/definition/11174/cold- sitehttps://www.techopedia.com/definition/29836/cloud-recovery

NEW QUESTION 234
A server administrator is configuring a new server that will hold large amounts of information. The server will need to be accessed by multiple users at the same time. Which of the following server roles will the administrator MOST likely need to install?

A. Messaging
B. Application
C. Print
D. Database

Answer: D

Explanation:
Few people are expected to use the database at the same time and users don’t need to customize the design of the database.
Reference:https://support.microsoft.com/en-us/office/ways-to-share-an-access-desktop-database-03822632-da43-4d8f-ba2a-68da245a0446
The server role that the administrator will most likely need to install for a server that will hold large amounts of information and will need to be accessed by multiple users at the same time is database. A database is a collection of structured data that can be stored, queried, manipulated, and analyzed using various methods and tools. A database server is a server that hosts one or more databases and provides access to them over a network. A database server can handle large amounts of information and support concurrent requests from multiple users or applications.

NEW QUESTION 237
A server administrator needs to implement load balancing without purchasing any new hardware or implementing any new software. Which of the following will the administrator most likely implement?

A. Round robin
B. Link aggregation
C. Most recently used
D. Heartbeat

Answer: B

Explanation:
Link aggregation is a technique that allows multiple network interfaces to act as one logical interface, increasing the bandwidth and redundancy of the connection. This can improve the load balancing of network traffic without requiring any new hardware or software. Round robin, most recently used, and heartbeat are not load balancing methods, but rather scheduling algorithms or monitoring techniques. References: CompTIA Server+ Certification Exam Objectives, Domain 2.0: Networking, Objective 2.3: Given a scenario, configure NIC teaming.

NEW QUESTION 239
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 243
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 246
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 249
An administrator is researching the upcoming licensing software requirements for an application that usually requires very little technical support. Which of the following licensing models would be the LOWEST cost solution?

A. Open-source
B. Per CPU socket
C. Per CPU core
D. Enterprise agreement

Answer: A

Explanation:
Open-source software is software that is freely available and can be modified and distributed by anyone. It usually requires very little technical support and has no licensing fees. Therefore, it would be the lowest cost solution for an application that does not need much support.References: https://www.comptia.org/training/resources/exam- objectives/comptia-server-sk0-005-exam-objectives (Objective 2.3)

NEW QUESTION 250
A server administrator needs to check remotely for unnecessary running services across 12 servers. Which of the following tools should the administrator use?

A. DLP
B. A port scanner
C. Anti-malware
D. A sniffer

Answer: B

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

NEW QUESTION 251
A server administrator purchased a single license key to use for all the new servers that will be imaged this year. Which of the following MOST likely refers to the licensing type that will be used?

A. Per socket
B. Open-source
C. Per concurrent user
D. Volume

Answer: D

Explanation:
This is the most likely licensing type that will be used because volume licensing allows a single license key to be used for multiple installations of a software product. Volume licensing is typically used by organizations that need to deploy software to a large number of devices or users.References:https://www.microsoft.com/en-us/licensing/licensing-programs/volume-licensing-programs

NEW QUESTION 255
A server technician installs a new NIC on a server and configures the NIC for IP connectivity. The technician then tests the connection using thepingcommand. Given the following partial output of thepingandipconfigcommands:

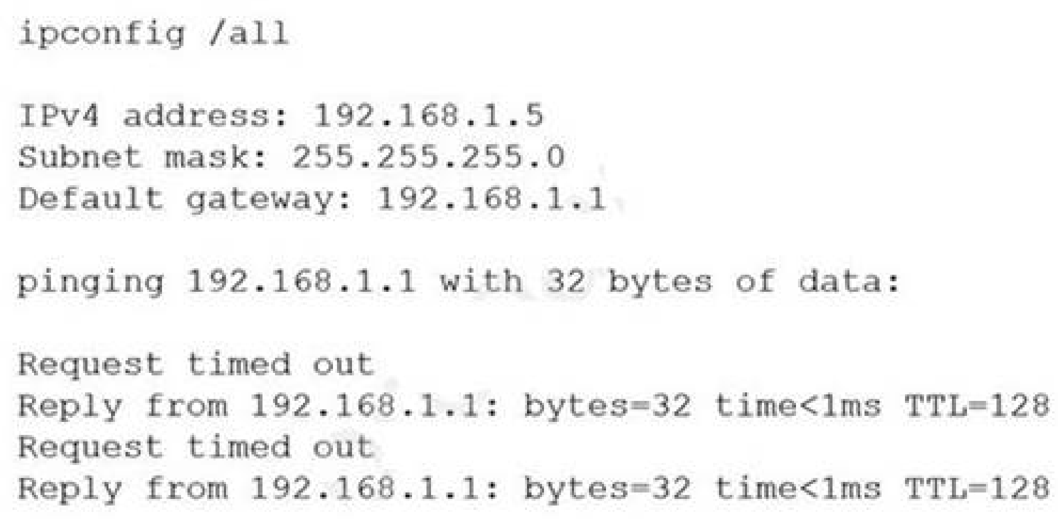

```
ipconfig /all

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 260
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 264
A technician is troubleshooting a server issue. The technician has determined several possible causes of the issue and has identified various solutions. Which of the following should the technician do next?

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

Answer: B

Explanation:
According to the CompTIA troubleshooting methodology, the fourth step is to establish a plan of action to resolve the problem and implement the solution1. The best practice is to test each solution individually to determine the root cause, rolling back the changes in between each test. This way, the technician can isolate the cause and avoid introducing new problems or making the situation worse. Testing each solution in succession and restoring the server from the latest snapshot (D) is not a good option because it may not identify the root cause and may overwrite important data. Implementing the shortest solution first to identify the issue and minimize downtime © is also not a good option because it may not solve the problem or may create new issues. Consulting internet forums to determine which is the most commoncause and deploy only that solution (A) is not a good option because it may not apply to the specific situation or may be outdated or inaccurate

NEW QUESTION 267
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 270
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 271
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 275
......