

CompTIA

Exam Questions SK0-005

CompTIA Server+ Certification Exam



NEW QUESTION 1

A snapshot is a feature that can be used in hypervisors to:

- A. roll back firmware updates.
- B. restore to a previous version.
- C. roll back application drivers.
- D. perform a backup restore.

Answer: B

Explanation:

A snapshot is a feature that can be used in hypervisors to restore to a previous version. A snapshot is a point-in-time copy of a virtual machine (VM) that captures the state and data of the VM at a specific moment. A snapshot can be created instantly and with minimal overhead, as it only stores the changes made to the VM after the snapshot was taken. A snapshot can be used to restore the VM to its previous state in case of data loss or corruption.

NEW QUESTION 2

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 3

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

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

Answer: A

Explanation:

The administrator should use a subnet mask of 255.255.255.0 for this setup. A subnet mask is a binary number that defines how many bits of an IP address are used for the network portion and how many bits are used for the host portion. The network portion identifies the specific network that the IP address belongs to, while the host portion identifies the specific device within that network. The subnet mask is usually written in dotted decimal notation, where each octet represents eight bits of the binary number. A 1 in the binary number means that the corresponding bit in the IP address is part of the network portion, while a 0 means that it is part of the host portion. For example, a subnet mask of 255.255.255.0 means that the first 24 bits (three octets) of the IP address are used for the network portion and the last 8 bits (one octet) are used for the host portion. This subnet mask allows up to 254 hosts per network ($2^8 - 2$). In this case, the IP address of 10.20.10.15 and the default gateway of 10.20.10.254 belong to the same network of 10.20.10.0/24 (where /24 indicates the number of bits used for the network portion), which can be defined by using a subnet mask of 255.255.255.0.

NEW QUESTION 4

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

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

Answer: D

Explanation:

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

NEW QUESTION 5

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

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

D. RAID 6

Answer: D

Explanation:

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

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

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

NEW QUESTION 6

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 7

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

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

Answer: D

Explanation:

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

NEW QUESTION 8

A server administrator is exporting Windows system files before patching and saving them to the following location:

\\server1\ITDept\

Which of the following is a storage protocol that the administrator is MOST likely using to save this data?

- A. eSATA
- B. FCoE
- C. CIFS
- D. SAS

Answer: C

Explanation:

The storage protocol that the administrator is most likely using to save data to the location \\server1\ITDept\ is CIFS. CIFS (Common Internet File System) is a protocol that allows file sharing and remote access over a network. CIFS is based on SMB (Server Message Block), which is a protocol that enables communication between devices on a network. CIFS uses UNC (Universal Naming Convention) paths to identify network resources, such as files or folders. A UNC path has the format \\servername\sharename\path\filename. In this case, server1 is the name of the server, ITDept is the name of the shared folder, and \ is the path within the shared folder.

NEW QUESTION 9

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 10

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 10

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 13

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 least amount of disk space lost to RAID overhead among the options, as it only uses one disk's worth of space for parity.

NEW QUESTION 14

DRAG DROP

A recent power Outage caused email services to go down. A server administrator also received alerts from the datacenter's UPS.

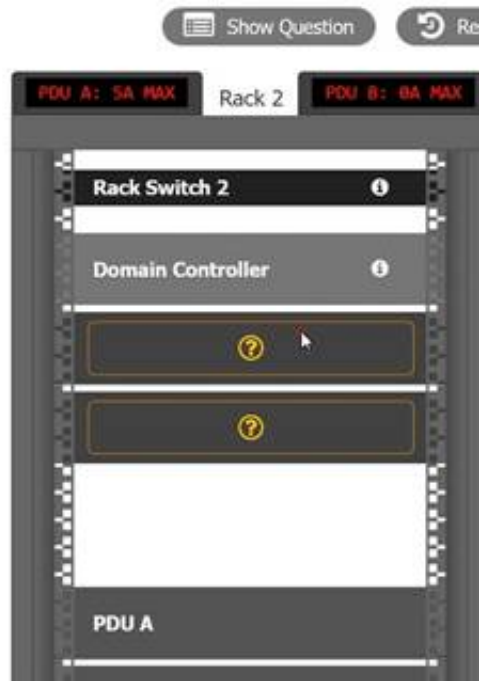
After some investigation, the server administrator learned that each POU was rated at a maximum Of 12A.

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

Data Center Racks 1 and 2



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Data Center Racks 1 and 2



NEW QUESTION 18

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 20

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

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

Answer: C

Explanation:

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

Reference:

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

NEW QUESTION 24

Which of the following script types would MOST likely be used on a modern Windows server OS?

- A. Batch
- B. VBS
- C. Bash
- D. PowerShell

Answer: D

Explanation:

PowerShell is a scripting language and a command-line shell that is designed for Windows server administration. It can perform various tasks such as configuration, automation, and management of servers and applications. Verified References: [PowerShell], [Scripting language]

NEW QUESTION 28

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

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

Answer: D

Explanation:

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

NEW QUESTION 32

Which of the following actions should a server administrator take once a new backup scheme has been configured?

- A. Overwrite the backups
- B. Clone the configuration
- C. Run a restore test
- D. Check the media integrity

Answer: C

Explanation:

The action that the server administrator should take once a new backup scheme has been configured is to run a restore test. A restore test is a process of verifying that the backup data can be successfully recovered and restored to its original location or a different location. A restore test can help ensure that the backup scheme is working properly, that the backup data is valid and consistent, and that there are no errors or issues during the recovery process. A restore test should be performed periodically and after any changes to the backup configuration or environment.

NEW QUESTION 36

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

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

Answer: C

Explanation:

The df -h command should be used to identify the issue of not being able to save large files to a directory on a Linux server. The df -h command displays disk space usage in human-readable format for all mounted file systems on the server. It shows the total size, used space, available space, percentage of use, and mount point of each file system. By using this command, a technician can check if there is enough free space on the file system where the directory is located or if it has reached its capacity limit.

NEW QUESTION 41

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

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

Answer: BF

Explanation:

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

NEW QUESTION 46

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 51

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 52

A server administrator is installing a new server with multiple NICs on it. The Chief Information Officer has asked the administrator to ensure the new server will have the least amount of network downtime but a good amount of network speed. Which of the following best describes what the administrator should implement on the new server?

- A. VLAN
- B. vNIC
- C. Link aggregation
- D. Failover

Answer: C

Explanation:

Link aggregation is the best option to implement on the new server to ensure the least amount of network downtime but a good amount of network speed. Link aggregation is a technique of combining multiple physical network interfaces into one logical interface to increase bandwidth, redundancy, and load balancing. Link aggregation can improve the performance and availability of the server by allowing it to use more than one network path for data transmission and failover in case of link failure. Link aggregation can be implemented using various protocols, such as IEEE 802.3ad (LACP), Cisco EtherChannel, or Linux bonding. References: [CompTIA Server+ Certification Exam Objectives], Domain 4.0: Networking, Objective 4.1: Given a scenario, configure network settings for servers.

NEW QUESTION 57

A technician wants to limit disk usage on a server. Which of the following should the technician implement?

- A. Formatting
- B. Compression
- C. Disk quotas
- D. Partitioning

Answer: C

Explanation:

Disk quotas are a way to limit disk usage on a server by setting a maximum amount of space that each user or group can use. Disk quotas can help manage disk space allocation, prevent disk space exhaustion, and enforce fair usage policies. Disk quotas can be set at the volume level or at the folder level, depending on the file system and operating system used. Reference: <https://docs.microsoft.com/en-us/windows-server/storage/ntfs/ntfs-disk-quotas-overview>

NEW QUESTION 60

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 65

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

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

Answer: BE

Explanation:

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

NEW QUESTION 68

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 71

A server technician is installing a new server OS on legacy server hardware. Which of the following should the technician do FIRST to ensure the OS will work as intended?

- A. Consult the HCL to ensure everything is supported.
- B. Migrate the physical server to a virtual server.
- C. Low-level format the hard drives to ensure there is no old data remaining.
- D. Make sure the case and the fans are free from dust to ensure proper cooling.

Answer: A

Explanation:

The first thing that the technician should do before installing a new server OS on legacy server hardware is to consult the HCL (Hardware Compatibility List) to ensure everything is supported. The HCL is a list of hardware devices and components that are tested and certified to work with a specific OS or software product. The HCL helps to avoid compatibility issues and performance problems that may arise from using unsupported or incompatible hardware. Migrating the physical server to a virtual server may be a good option to improve scalability and flexibility, but it requires additional hardware and software resources and may not be feasible for legacy server hardware. Low-level formatting the hard drives may be a good practice to erase any old data and prepare the drives for a new OS installation, but it does not guarantee that the hardware will work with the new OS. Making sure the case and the fans are free from dust may be a good practice to ensure proper cooling and prevent overheating, but it does not guarantee that the hardware will work with the new OS. References: <https://www.howtogeek.com/190014/virtualization-basics-understanding-techniques-and-fundamentals/> <https://www.howtogeek.com/173353/how-to-low-level-format-or-write-zeros-to-a-hard-drive/> <https://www.howtogeek.com/303282/how-to-manage-your-pcs-fans-for-optimal-airflow-and-cooling/>

NEW QUESTION 75

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

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

Answer: B

Explanation:

The port that must remain open for a secure website is port 443. Port 443 is used by Hypertext Transfer Protocol Secure (HTTPS), which is an extension of HTTP that encrypts and authenticates the communication between a web server and a web browser. HTTPS ensures that the data transmitted over the web is protected from eavesdropping, tampering, or spoofing. Therefore, port 443 must remain open for a secure website to function properly. References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 2, Lesson 2.2, Objective 2.2

NEW QUESTION 80

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 82

The accounting department needs more storage and wants to retain the current data for quick read-write access. The accounting server does not have any internet drive bays available to keep both disks however the server does have USB 3.0 and eSATA ports available. Which of the following is the BEST way to accomplish the department's goals?

- A. Copy the existing data to an external USB 3.0 enclosure.
- B. Place the existing data on a DVD and use the internal DVD-ROM drive.
- C. Transfer the existing data to an external eSATA enclosure.
- D. Move the existing data to a new
- E. larger internal hard drive.

Answer: C

Explanation:

The best way to accomplish the department's goals is to transfer the existing data to an external eSATA enclosure, which is a device that connects an external hard drive to a computer using an eSATA port. This will allow the accounting department to retain the current data for quick read-write access, as eSATA provides high-speed data transfer rates and supports hot-plugging. Unlike USB 3.0, eSATA does not share bandwidth with other devices, which can improve performance and reliability. Verified References: [eSATA vs USB 3.0]

NEW QUESTION 86

A server technician has received reports of database update errors. The technician checks the server logs and determines the database is experiencing synchronization errors. To attempt to correct the errors, the technician should FIRST ensure:

- A. the correct firewall zone is active
- B. the latest firmware was applied
- C. NTP is running on the database system
- D. the correct dependencies are installed

Answer: C

Explanation:

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

NEW QUESTION 89

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

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

Answer: C

Explanation:

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

NEW QUESTION 91

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

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

Answer: B

Explanation:

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

NEW QUESTION 93

A server administrator is completing an OS installation for a new server. The administrator patches the server with the latest vendor-suggested software, configures DHCP, and verifies all network cables are properly connected in the IDF, but there is no network connectivity. Which of the following is the MOST likely reason for the lack of connectivity?

- A. The VLAN is improperly configured.
- B. The DNS configuration is invalid.
- C. The OS version is not compatible with the network switch vendor.
- D. The HIDS is preventing the connection.

Answer: A

Explanation:

If the server administrator patches the server with the latest vendor-suggested software, configures DHCP, and verifies all network cables are properly connected in the IDF, but there is no network connectivity, then the most likely reason for the lack of connectivity is that the VLAN is improperly configured. A VLAN (Virtual Local Area Network) is a logical grouping of network devices that share the same broadcast domain and can communicate with each other without routing. If the server is assigned to a different VLAN than the DHCP server or the default gateway, it will not be able to obtain an IP address or reach other network devices. The DNS configuration is not relevant for network connectivity, as DNS only resolves names to IP addresses. The OS version is not likely to be incompatible with the network switch vendor, as most network switches use standard protocols and interfaces. The HIDS (Host-based Intrusion Detection System) is not likely to prevent the connection, as HIDS only monitors and alerts on suspicious activities on the host. References: <https://www.howtogeek.com/190014/virtualization-basics-understanding-techniques-and-fundamentals/> <https://www.howtogeek.com/164981/how-to-use-nslookup-to-check-domain-name-information-in-microsoft-windows/> <https://www.howtogeek.com/202794/what-is-an-intrusion-detection-system-ids-and-how-does-it-work/>

NEW QUESTION 95

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 100

An organization implements split encryption keys for sensitive files. Which of the following types of risks does this mitigate?

- A. Hardware failure
- B. Malware
- C. Data corruption
- D. Insider threat

Answer: D

Explanation:

An insider threat is a type of risk that can be mitigated by implementing split encryption keys for sensitive files. 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. An insider threat can cause data breaches, sabotage, fraud, theft, espionage, or other damages to the organization. Split encryption keys are a method of encrypting data using multiple keys that are stored separately and require collaboration to decrypt. Split encryption keys can prevent an insider threat from accessing or compromising sensitive data without being detected by another authorized party who holds another key. Hardware failure is a type of risk that involves physical damage or malfunction of hardware components such as hard drives, memory modules, power supplies, or fans. Hardware failure can cause data loss, system downtime, performance issues, or other problems for the organization. Hardware failure cannot be mitigated by split encryption keys, but by backup, redundancy, monitoring, and maintenance measures.

NEW QUESTION 105

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

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

Answer: B

Explanation:

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

? Lower cost: Cloud IaaS can reduce the capital and operational expenses of

building and maintaining a physical disaster recovery site, as users only pay for the resources they use on demand¹².

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

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

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

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

? Simplicity: Cloud IaaS can simplify the failover process, as users can use the cloud provider's tools and services to automate and orchestrate the failover operations¹².

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

NEW QUESTION 108

A security manager is concerned that a rogue employee could boot a server from an outside USB drive. Which of the following actions can be taken to reduce this risk? (Select TWO).

- A. Close unneeded ports.
- B. Disable unneeded physical ports.
- C. Set a BIOS password.

- D. Install a SIEM.
- E. Disable unneeded services.
- F. Install a HIDS.

Answer: BC

Explanation:

Disabling unneeded physical ports would prevent unauthorized devices from being connected to the server, such as an outside USB drive. Setting a BIOS password would restrict access to the boot settings and prevent unauthorized changes to the boot order. The other options would not address the risk of booting from an outside USB drive

NEW QUESTION 111

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: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 5, Lesson 5.2, Objective 5.2

NEW QUESTION 116

Network connectivity to a server was lost when it was pulled from the rack during maintenance. Which of the following should the server administrator use to prevent this situation in the future?

- A. Cable management
- B. Rail kits
- C. A wireless connection
- D. A power distribution unit

Answer: A

Explanation:

The server administrator should use cable management to prevent network connectivity loss when pulling a server from the rack during maintenance. Cable management is a practice of organizing and securing the cables that connect various devices and components in a system. Cable management can help improve airflow, reduce clutter, prevent tangling, and avoid accidental disconnection or damage of cables. Cable management can be done using various tools and techniques, such as cable ties, cable trays, cable labels, cable organizers, or cable ducts.

NEW QUESTION 120

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 124

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

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

Answer: D

Explanation:

Grandfather-father-son (GFS) is a common backup rotation scheme that uses daily, weekly, and monthly backup media in a first-in, first-out (FIFO) method. The daily backups are rotated on a 3-months basis using a FIFO system as above. The weekly backups are similarly rotated on a bi-yearly basis, and the monthly backups are rotated on an annual basis. The oldest backup media in each cycle are overwritten by the newest ones. This scheme provides multiple versions of backup data at different intervals, allowing for flexible restoration options. Waterfall is another name for GFS. Synthetic full is a backup method that combines an

initial full backup with subsequent incremental backups to create a new full backup without transferring all data again. Tower of Hanoi is another backup rotation scheme that uses an algorithm based on moving disks between three pegs. References:
? https://en.wikipedia.org/wiki/Backup_rotation_scheme

NEW QUESTION 126

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 128

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 131

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

- A. 1
- B. 5
- C. 6

Answer: D

Explanation:

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

? https://en.wikipedia.org/wiki/Standard_RAID_levels

NEW QUESTION 135

Which of the following is an example of load balancing?

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

D. Failover

Answer: A

Explanation:

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

NEW QUESTION 140

A technician installed a kernel upgrade to a Linux server. The server now crashes regularly. Which of the following is the most likely cause?

- A. Necessary dependencies were installed for multiple architectures.
- B. There is not enough hard drive space.
- C. The server is infected with a virus.
- D. Some modules are not compatible.

Answer: D

Explanation:

A kernel upgrade is a process of updating the core component of a Linux operating system that manages the hardware, memory, processes, and drivers. A kernel upgrade can improve the performance, security, and compatibility of the system, but it can also introduce errors if some modules are not compatible with the new kernel version. Modules are pieces of code that can be loaded and unloaded into the kernel to provide additional functionality or support for specific devices. If a module is not compatible with the kernel, it can cause crashes or instability.

References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 4, Lesson 4.2, Objective 4.2

NEW QUESTION 143

A company's security team has noticed employees seem to be blocking the door in the main data center when they are working on equipment to avoid having to gain access each time. Which of the following should be implemented to force the employees to enter the data center properly?

- A. A security camera
- B. A mantrap
- C. A security guard
- D. A proximity card

Answer: B

Explanation:

A mantrap is a security device that consists of two interlocking doors that allow only one person to enter at a time. A mantrap would prevent employees from blocking the door in the main data center and force them to enter properly using their credentials. The other options would not enforce proper entry to the data center

NEW QUESTION 145

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

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

Answer: C

Explanation:

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

References: https://en.wikipedia.org/wiki/Standard_RAID_levels

NEW QUESTION 150

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

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

Answer: BE

Explanation:

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

References:

CompTIA Server+ Certification Exam Objectives1, page 12

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

NEW QUESTION 151

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 153

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 155

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 159

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

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

Answer: A

Explanation:

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

NEW QUESTION 163

A company needs to increase the security controls on its servers. An administrator is implementing MFA on all servers using cost effective techniques. Which of the following should the administrator use to satisfy the MFA requirement?

- A. Biometrics
- B. Push notifications
- C. Smart cards
- D. Physical tokens

Answer: B

Explanation:

Push notifications are messages that are sent from an application or a service to a user's device without requiring the user to open or request them. They can be used as a cost-effective technique for implementing MFA (Multi-Factor Authentication) on servers by sending verification codes or approval requests to the user's smartphone or tablet when they try to log in to the server. Verified References: [Push notifications], [MFA]

NEW QUESTION 167

A technician set up a new multifunction printer. After adding the printer to the print server, the technician configured the printer on each user's machine. Several days later, users reported that they were no longer able to print, but scanning to email worked. Which of the following is most likely causing this issue?

- A. The gateway is no longer being reached.
- B. The network firewall was enabled.
- C. The printer's network interface failed.
- D. The printer had DHCP enabled.

Answer: D

Explanation:

The most likely cause of this issue is that the printer had DHCP enabled, which changed its IP address after adding it to the print server and configuring it on each user's machine. DHCP (Dynamic Host Configuration Protocol) is a network protocol that assigns IP addresses and other network configuration parameters to devices automatically, without manual intervention. DHCP can simplify network management and avoid IP conflicts, but it can also cause problems if the devices are not configured to use static or reserved IP addresses. If the printer had DHCP enabled, it might have received a different IP address from the DHCP server after rebooting or reconnecting to the network, which would make it unreachable by the print server and the users' machines that were configured with the previous IP address. Scanning to email would still work, as it does not depend on the print server or the users' machines, but on the printer's SMTP settings and internet connection. References: [CompTIA Server+ Certification Exam Objectives], Domain 4.0: Networking, Objective 4.1: Given a scenario, configure network settings for servers.

NEW QUESTION 169

An administrator notices high traffic on a certain subnet and would like 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 170

A technician has received multiple reports of issues with a server. The server occasionally has a BSOD, powers off unexpectedly, and has fans that run continuously. Which of the following BEST represents what the technician should investigate during troubleshooting?

- A. Firmware incompatibility
- B. CPU overheating
- C. LED indicators
- D. ESD issues

Answer: B

Explanation:

Unexpected shutdowns. If the system is randomly shutting down or rebooting, the most likely cause is a heat problem.
Reference: <https://www.microsoftpressstore.com/articles/article.aspx?p=2224043&seqNum=3>

NEW QUESTION 173

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

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

D. #

Answer: D

Explanation:

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

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

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

NEW QUESTION 177

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

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

Answer: D

Explanation:

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

NEW QUESTION 182

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 of numerical 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 184

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

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

Answer: C

Explanation:

A synthetic full backup is a type of backup that copies changed data from a server and then combines the backups on the backup target. This way, the backup target always has a full backup of the server, without requiring a full backup to be performed over the network. A synthetic full backup reduces the network bandwidth and time required for backups, while also simplifying the restoration process¹

NEW QUESTION 185

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 186

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

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

Answer: A

Explanation:

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

NEW QUESTION 190

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 192

Which of the following cloud models is BEST described as running workloads on resources that are owned by the company and hosted in a company-owned data center, as well as on rented servers in another company's data center?

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

Answer: B

Explanation:

This is the best description of a hybrid cloud model because it combines both private and public cloud resources. A private cloud is a cloud environment that is owned and operated by a single organization and hosted in its own data center. A public cloud is a cloud environment that is owned and operated by a third-party provider and hosted in its data center. A hybrid cloud allows an organization to leverage both types of cloud resources depending on its needs and preferences. References: <https://azure.microsoft.com/en-us/overview/what-is-hybrid-cloud-computing/>

NEW QUESTION 197

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

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

Answer: D

Explanation:

The first thing that a technician should verify before decommissioning and wiping a file server is non-utilization, which means that no one is using or accessing the server or its data. This can be done by checking logs, monitoring network traffic, or contacting users or stakeholders. Non-utilization ensures that decommissioning

and wiping will not cause any data loss or disruption to business operations. Verified References: [Server Decommissioning Checklist]

NEW QUESTION 199

Which of the following is an architectural reinforcement that attempts to conceal the interior of an organization?

- A. Bollards
- B. Signal blocking
- C. Reflective glass
- D. Data center camouflage

Answer: C

Explanation:

Reflective glass is an architectural reinforcement that attempts to conceal the interior of an organization by reflecting light and preventing outsiders from seeing inside. Reflective glass can also reduce heat and glare, and enhance the aesthetic appearance of a building. Reflective glass is often used in high-security facilities, such as data centers, government buildings, or corporate headquarters¹²

1: Server Architecture for CompTIA Server+ (SK0-004) | Pluralsight 2: Introducing the CompTIA Infrastructure Career Pathway

NEW QUESTION 201

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 204

A VLAN needs to be configured within a virtual environment for a new VM. Which of the following will ensure the VM receives a correct IP address?

- A. A virtual router
- B. A host NIC
- C. A VPN
- D. A virtual switch
- E. A vNIC

Answer: D

Explanation:

The correct answer is D. A virtual switch.

A virtual switch is a software-based network device that connects the virtual machines (VMs) in a virtual environment and allows them to communicate with each other and with the physical network. A virtual switch can also create and manage virtual LANs (VLANs), which are logical segments of a network that separate the traffic of different VMs or groups of VMs. A VLAN needs a DHCP server to assign IP addresses to the VMs that belong to it. A virtual switch can act as a DHCP relay agent and forward the DHCP requests from the VMs to the DHCP server on the physical network. This way, the VMs can receive correct IP addresses for their VLANs¹²³

A virtual router is a software-based network device that routes packets between different networks or subnets. A virtual router can also create and manage VLANs, but it is not necessary for a VM to receive a correct IP address. A virtual router can be used to provide additional security, redundancy, or load balancing for the VMs¹²

A host NIC is a physical network interface card that connects the host machine to the physical network. A host NIC can also support VLAN tagging, which allows the host machine to communicate with different VLANs on the network. However, a host NIC alone cannot ensure that a VM receives a correct IP address for its VLAN. The host NIC needs to be connected to a virtual switch that can relay the DHCP requests from the VMs to the DHCP server¹²

A VPN is a virtual private network that creates a secure tunnel between two or more devices over the internet. A VPN can be used to encrypt and protect the data traffic of the VMs, but it is not related to the configuration of VLANs or IP addresses. A VPN does not affect how a VM receives a correct IP address for its VLAN¹⁴

A vNIC is a virtual network interface card that connects a VM to a virtual switch or a virtual router. A vNIC can also support VLAN tagging, which allows the VM to communicate with different VLANs on the network. However, a vNIC alone cannot ensure that a VM receives a correct IP address for its VLAN. The vNIC needs to be connected to a virtual switch or a virtual router that can relay the DHCP requests from the VMs to the DHCP server¹²

NEW QUESTION 208

Which of the following commands should a systems administrator use to create a batch script to map multiple shares'?

- A. `nbtstat`
- B. `netuse`
- C. `tracert`

D. netstst

Answer: B

Explanation:

The net use command is a Windows command that can be used to create a batch script to map multiple shares. The net use command can connect or disconnect a computer from a shared resource, such as a network drive or a printer, or display information about computer connections. The syntax of the net use command is:

```
net use [devicename | *] [\computername\sharename[\u0003volume] [password | *]] [/user:[domainname\]username] [/user:[dotted domain name\]username] [/user:[username@dotted domain name] [/savecred] [/smartcard] [{/delete | /persistent:{yes | no}}] where:
```

devicename = the drive letter or printer port to assign to the shared resource computername = the name of the computer that provides access to the shared resource sharename = the name of the shared resource password = the password needed to access the shared resource /user = specifies a different username to make the connection

/savecred = stores the provided credentials for future use /smartcard = uses a smart card for authentication /delete = cancels a network connection and removes the connection from the list of persistent connections /persistent = controls whether the connection is restored at logon

To create a batch script to map multiple shares, you can use the net use command with different drive letters and share names, for example:

```
net use W: \\computer1\share1 net use X: \\computer2\share2 net use Y: \\computer3\share3
```

You can also add other options, such as passwords, usernames, or persistence, as needed. To save the batch script, you can use Notepad or any text editor and save the file with a .bat extension12.

References: 1 <https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/net-use> 2 <https://www.watchingthenet.com/create-a-batch-file-to-map-drives-folders.html>

NEW QUESTION 211

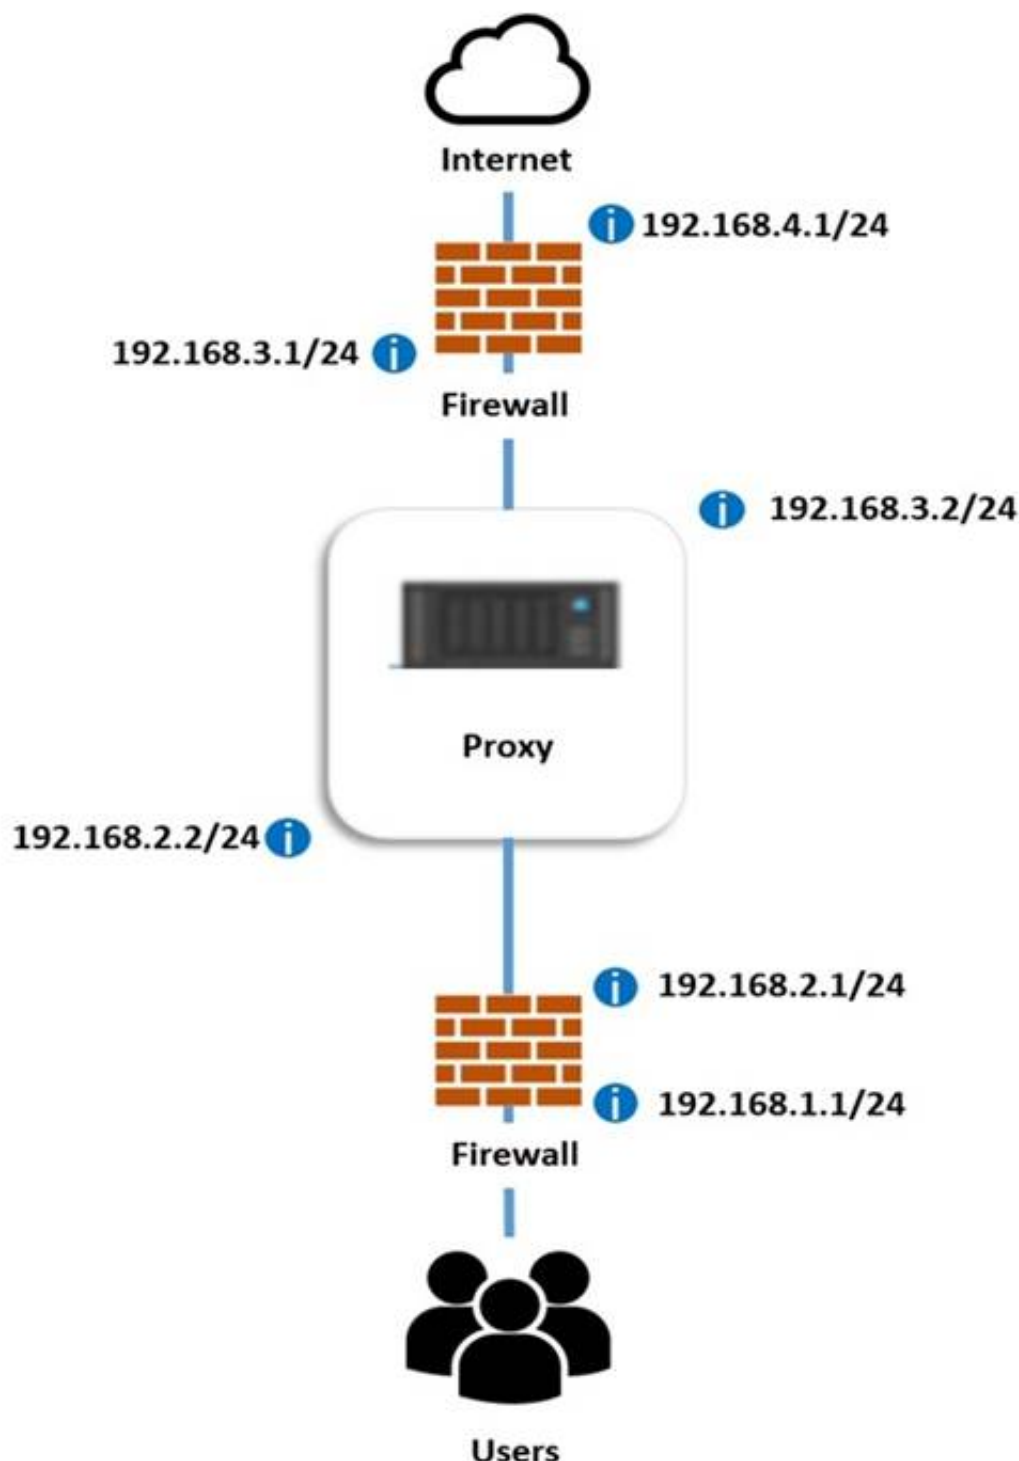
HOTSPOT

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	192.168.4.1
		192.168.4.0	192.168.1.1
		192.168.1.1	192.168.3.0
		192.168.2.0	192.168.1.0
		192.168.1.0	192.168.2.2
		192.168.4.1	0.0.0.0
		192.168.2.1	192.168.3.1
		0.0.0.0	255.255.255.0
		192.168.3.1	192.168.3.2
		255.255.255.0	192.168.4.0
		192.168.3.2	192.168.2.1
		192.168.2.2	192.168.2.0
192.168.1.0	255.255.255.0	▼	▼
		192.168.3.0	192.168.4.1
		192.168.4.0	192.168.1.1
		192.168.1.1	192.168.3.0
		192.168.2.0	192.168.1.0
		192.168.1.0	192.168.2.2
		192.168.4.1	0.0.0.0
		192.168.2.1	192.168.3.1
		0.0.0.0	255.255.255.0
		192.168.3.1	192.168.3.2
		255.255.255.0	192.168.4.0
		192.168.3.2	192.168.2.1
		192.168.2.2	192.168.2.0

- 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	192.168.4.1
		192.168.4.0	192.168.1.1
		192.168.1.1	192.168.3.0
		192.168.2.0	192.168.1.0
		192.168.1.0	192.168.2.2
		192.168.4.1	0.0.0.0
		192.168.2.1	192.168.3.1
		0.0.0.0	255.255.255.0
		192.168.3.1	192.168.3.2
		255.255.255.0	192.168.4.0
		192.168.3.2	192.168.2.1
		192.168.2.2	192.168.2.0
192.168.1.0	255.255.255.0	▼	▼
		192.168.3.0	192.168.4.1
		192.168.4.0	192.168.1.1
		192.168.1.1	192.168.3.0
		192.168.2.0	192.168.1.0
		192.168.1.0	192.168.2.2
		192.168.4.1	0.0.0.0
		192.168.2.1	192.168.3.1
		0.0.0.0	255.255.255.0
		192.168.3.1	192.168.3.2
		255.255.255.0	192.168.4.0
		192.168.3.2	192.168.2.1
		192.168.2.2	192.168.2.0

NEW QUESTION 216

A server technician is installing application updates on a Linux server. When the technician tries to install a MySQL update, the GUI displays the following error message: AVC denial. Which of the following should the technician do for the MySQL update to install?

- A. Download the update manually and run a checksum utility to verify file integrity.
 B. Issue the setenforce 0 command.
 C. Create a firewall rule to allow port 3306 through the firewall.
 D. Issue the yum -y update mysql command.

Answer: B

Explanation:

The AVC denial error message indicates that SELinux (Security-Enhanced Linux) is preventing the MySQL update from installing. SELinux is a security module that enforces mandatory access control policies on Linux systems. To install the MySQL update, the technician should issue the `setenforce 0` command, which temporarily disables SELinux enforcement until the next reboot. Downloading the update manually, creating a firewall rule, or issuing the `yum -y update mysql` command will not resolve the error. References: [CompTIA Server+ Certification Exam Objectives], Domain 4.0: Server Administration, Objective 4.3: Given a scenario, troubleshoot server issues using appropriate tools.

NEW QUESTION 218

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 servers.

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 conflicts.

NEW QUESTION 220

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

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

Answer: D

Explanation:

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

NEW QUESTION 222

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

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

Answer: C

Explanation:

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

NEW QUESTION 223

A data center has 4U rack servers that need to be replaced using VMs but without losing any data. Which of the following methods will MOST likely be 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 226

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

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

Answer: B

Explanation:

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

NEW QUESTION 227

A server administrator has received tickets from users who report the system runs very slowly and various unrelated messages pop up when they try to access an internet-facing web application using default ports. The administrator performs a scan to check for open ports and reviews the following report:

Starting Nmap 7.70 <https://nmap.org>) at 2019-09-19 14:30 UTC Nmap scan report for www.abc.com (172.45.6.85)

Host is up (0.0021s latency)

Other addresses for www.abc.com (not scanned) : 4503 : F7b0 : 4293: 703: : 3209 RDNS record for 172.45.6.85: 1ga45s12-in-f1.2d100.net

Port State Service 21/tcp filtered ftp 22/tcp filtered ssh 23/tcp filtered telnet

69/tcp open @username.com 80/tcp open http

110/tcp filtered pop 143/tcp filtered imap 443/tcp open https

1010/tcp open www.popup.com 3389/tcp filtered ms-abc-server

Which of the following actions should the server administrator perform on the server?

- A. Close ports 69 and 1010 and rerun the scan.
- B. Close ports 80 and 443 and rerun the scan.
- C. Close port 3389 and rerun the scan.
- D. Close all ports and rerun the scan.

Answer: A

Explanation:

Port 69 is used for TFTP (Trivial File Transfer Protocol), which is an insecure and unencrypted protocol for file transfer. Port 1010 is used for a malicious website that generates pop-up ads. Both of these ports are likely to be exploited by hackers or malware to compromise the server or the web application. The server administrator should close these ports and rerun the scan to verify that they are no longer open.

References = 1: Why Are Some Network Ports Risky, And How Do You Secure Them? - How-To Geek(<https://www.howtogeek.com/devops/why-are-some-ports-risky-and-how-do-you-secure-them/>) 2: Switchport Port Security Explained With Examples -

ComputerNetworkingNotes(<https://www.computernetworkingnotes.com/ccna-study-guide/switchport-port-security-explained-with-examples.html>)

NEW QUESTION 232

A Linux server was recently updated. Now, the server stops during the boot process with a blank screen and an f prompt. Which 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. UEFI (Unified Extensible Firmware Interface) is a standard that defines the interface and functionality of the firmware that initializes the hardware and software components of a system before loading the operating system. UEFI boot is a process that uses UEFI firmware to load and execute a boot loader, which is a program that loads the operating system kernel and other essential files. A Linux boot file is a file that contains information and instructions for the boot loader, such as the location of the kernel, the root file system, and the boot parameters. If a Linux boot file is missing or corrupted, the boot loader cannot find or load the kernel, and the system stops during the boot process with a blank screen and an f prompt.

References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 4, Lesson 4.1, Objective 4.1

NEW QUESTION 236

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 239

A server administrator has a system requirement to install the virtual OS on bare metal hardware. Which of the following hypervisor virtualization technologies should the administrator use to BEST meet the system requirements? (Select TWO)

- A. Host
- B. Template
- C. Clone
- D. Type1
- E. Type2
- F. Guest

Answer: BD

Explanation:

A template is a preconfigured virtual machine image that can be used to create new virtual machines quickly and easily. A template can include the operating system, applications, settings, and data that are required for a specific purpose or role. A type 1 hypervisor is a virtualization technology that runs directly on bare metal hardware, without requiring an underlying operating system. A type 1 hypervisor can provide better performance, security, and isolation for virtual machines than a type 2 hypervisor, which runs on top of an operating system. Verified References: [Template], [Type 1 hypervisor]

NEW QUESTION 240

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

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

Answer: B

Explanation:

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

NEW QUESTION 242

The HIDS logs on a server indicate a significant number of unauthorized access attempts via USB devices at startup. Which of the following steps should a server administrator take to BEST secure the server without limiting functionality?

- A. Set a BIOS/UEFI password on the server.
- B. Change the boot order on the server and restrict console access
- C. Configure the host OS to deny login attempts via USB.
- D. Disable all the USB ports on the server.

Answer: B

Explanation:

Changing the boot order on the server and restricting console access would prevent unauthorized access attempts via USB devices at startup, as the server would not boot from any external media and only authorized users could access the console. Setting a BIOS/UEFI password on the server would also help, but it could be bypassed by resetting the CMOS battery or using a backdoor password. Configuring the host OS to deny login attempts via USB would not prevent booting from a malicious USB device that could compromise the system before the OS loads. Disabling all the USB ports on the server would limit functionality, as some peripherals or devices may need to use them. References:

- ? <https://www.pcmag.com/how-to/dont-plug-it-in-how-to-prevent-a-usb-attack>
- ? <https://www.techopedia.com/definition/10362/boot-order>
- ? <https://www.techopedia.com/definition/10361/console-access>
- ? <https://www.techopedia.com/definition/102/bios-password>
- ? <https://www.techopedia.com/definition/10363/cmos-battery>

NEW QUESTION 243

A technician has installed a new 4TB hard drive in a Windows server. Which of the following should the technician perform FIRST to provision the new drive?

- 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 GPT.

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 247

A technician is able to copy a file 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 249

A newly installed server is accessible to local users, but remote users are unable to connect. Which of the following is MOST likely misconfigured?

- A. The IP address
- B. The default gateway
- C. The VLAN
- D. The subnet mask

Answer: B

Explanation:

This is the most likely misconfigured setting because the default gateway is the router that connects the local network to other networks. If the default gateway is incorrect, the server will not be able to communicate with remote users or devices outside its own subnet. References: <https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/ipconfig>

NEW QUESTION 253

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

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

Answer: A

Explanation:

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

NEW QUESTION 254

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-site> <https://www.techopedia.com/definition/29836/cloud-recovery>

NEW QUESTION 257

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 260

An administrator has been troubleshooting a server issue. The administrator carefully questioned the users and examined the available logs. Using this information, the administrator was able to rule out several possible causes and develop a theory as to what the issue might be. Through further testing, the administrator's theory proved to be correct. Which of the following should be the next step to troubleshoot the issue?

- A. Document the findings and actions.
- B. Escalate the issue to the management team.
- C. Implement the solution.
- D. Establish an action plan.

Answer: D

Explanation:

The next step to troubleshoot the issue after developing and testing a theory is to establish an action plan. This involves identifying the steps needed to implement the solution, estimating the time and resources required, and evaluating the potential risks and impacts of the solution. Documenting the findings and actions, escalating the issue to the management team, or implementing the solution are steps that should be done after establishing an action plan. References: [CompTIA Server+ Certification Exam Objectives], Domain 6.0: Disaster Recovery, Objective 6.2: Explain troubleshooting theory and methodologies.

NEW QUESTION 264

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

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

Answer: AC

Explanation:

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

Reference:

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

NEW QUESTION 267

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 268

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

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

Answer: CF

Explanation:

The best immediate actions to prevent unauthorized server access are to stop all unneeded services and block the ports on the firewall. Stopping unneeded services reduces the attack surface of the server by eliminating potential entry points for attackers. For example, if the server is not functioning as a web server, there is no need to run a web service on port 80. Blocking ports on the firewall prevents unauthorized network traffic from reaching the server. For example, if port 80 is not needed for any legitimate purpose, it can be blocked on the firewall to deny any connection attempts on that port.

NEW QUESTION 270

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

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

Answer: A

Explanation:

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

NEW QUESTION 273

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 277

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 281

A systems administrator needs to back up changes made to a data store on a daily basis during a short time frame. The administrator wants to maximize RTO when restoring data. Which of the following backup methodologies would best fit this scenario?

- A. Off-site backups
- B. Full backups
- C. Differential backups
- D. Incremental backups

Answer: D

Explanation:

An incremental backup is a backup method that only backs up the files that have changed since the last backup, whether it was a full or an incremental backup. An incremental backup can save disk space and time, as it only copies the new or modified data. An incremental backup can also improve the RTO (Recovery Time Objective), which is the maximum acceptable time to restore data after a disaster. This is because an incremental backup can restore data faster than a full or a differential backup, as it only needs to apply the latest changes to the previous backup¹.

NEW QUESTION 284

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 285

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 common cause and deploy only that solution (A) is not a good option because it may not apply to the specific situation or may be outdated or inaccurate

NEW QUESTION 288

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 290

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 293

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SK0-005 Practice Exam Features:

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

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