

Exam Questions SK0-005

CompTIA Server+ Certification Exam

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



NEW QUESTION 1

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 2

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

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

Answer: C

Explanation:

The administrator should use the command `mount -o remount,rw /newdisk` to successfully mount the drive with the required parameters. The mount command is used to mount file systems on Linux systems. The `-o` option specifies options for mounting file systems. The `remount` option re-mounts an already mounted file system with different options. The `rw` option mounts a file system with read-write permissions. In this case, /newdisk is a mount point for a new hard disk that was mounted with read-only permissions by default. To allow users to create directories or files on /newdisk, the administrator needs to re-mount /

Reference:

<https://unix.stackexchange.com/QUESTION NO:s/149399/how-to-remount-as-read-write-a-specific-mount-of-device>

NEW QUESTION 3

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

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

Answer: D

Explanation:

RAID 6 is a type of RAID level that uses two parity blocks to provide fault tolerance and redundancy for data storage. RAID 6 can withstand the failure of up to two disks in the array without losing any data. RAID 6 requires a minimum of four disks to operate, and it distributes the data and parity blocks across all the disks in the array. RAID 6 has a high write penalty, which means that it takes more time and resources to write data to the disks than to read data from them. However, RAID 6 offers a 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 4

An administrator is rebooting servers manually after a group of updates were deployed through SCCM. The administrator notices several of the servers did not receive the deployed update. Which of the following should the administrator review first?

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

Answer: D

Explanation:

The first thing the administrator should check is whether the server is in the collection that was scheduled to receive the update through SCCM. A collection is a

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

NEW QUESTION 5

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 6

A systems administrator recently upgraded the memory in a server, and now the server does not turn on, and nothing is displayed on the screen. Which of the following is the next step the administrator should take to diagnose the error without opening the machine?

- A. Perform a cold reboot.
- B. Listen for POST code beeps.
- C. Call technical support.
- D. Check the monitor connection.

Answer: B

Explanation:

A power-on self-test (POST) is a diagnostic process that runs when a server is turned on to check the basic functionality of the hardware components and report any errors or faults. A POST code is a series of beeps or flashes that indicate the status of the POST process and identify any problems that prevent the server from booting up. A POST code can be heard through a speaker or seen on a display attached to the server motherboard. A POST code is useful for diagnosing errors without opening the machine or using any software tools.

References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 5, Lesson 5.1, Objective 5.1

NEW QUESTION 7

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 8

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

- A. Tabletop
- B. Backup recovery test
- C. Live failover
- D. Hot-site visit audit

Answer: A

Explanation:

A tabletop exercise is a type of disaster recovery testing exercise that is primarily used to discuss incident response strategies for critical systems without affecting production data. A tabletop exercise is a discussion-based session where team members meet in an informal, classroom setting to review their roles and responsibilities during an emergency and their responses to a hypothetical scenario. A facilitator guides the participants through the discussion and evaluates the strengths and weaknesses of the preparedness program. A tabletop exercise does not involve any actual deployment of resources or activation of systems. A backup recovery test (B) is a type of disaster recovery testing exercise that involves restoring data from backup media to verify its integrity and availability. A

backup recovery test may affect production data if it is not performed on a separate environment. A live failover © is a type of disaster recovery testing exercise that involves switching operations from a primary site to a secondary site in case of a failure or disruption. A live failover may affect production data if it is not performed on a simulated environment. A hot-site visit audit (D) is a type of disaster recovery testing exercise that involves inspecting and evaluating a hot site, which is a backup location that has fully operational equipment and resources to resume business operations in case of a disaster. A hot-site visit audit does not involve any discussion of incident response strategies or simulation of scenarios. References: 1
<https://www.ready.gov/testing-exercises2> <https://www.ready.gov/exercises>

NEW QUESTION 9

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

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

Answer: C

Explanation:

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

NEW QUESTION 10

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

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

Answer: B

Explanation:

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

NEW QUESTION 10

A technician noted the RAID hard drives were functional while troubleshooting a motherboard failure. The technician installed a spare motherboard with similar specifications and used the original components. Which of the following should the technician do to restore operations with minimal downtime?

- A. Reinstall the OS and programs.
- B. Configure old drives to RAID.
- C. Reconfigure the RAID.
- D. Install from backup.

Answer: C

Explanation:

RAID (Redundant Array of Independent Disks) is a technology that combines multiple hard drives into a logical unit that provides improved performance, reliability, or capacity. RAID can be implemented by hardware, software, or a combination of both. Hardware RAID uses a dedicated controller to manage the RAID array, while software RAID uses the operating system or a driver to do the same¹. In this scenario, the technician noted that the RAID hard drives were functional while troubleshooting a motherboard failure. This means that the data on the drives was not corrupted or lost. However, the technician installed a spare motherboard with similar specifications and used the original components. This means that the new motherboard may not have the same RAID configuration as the old one, or it may not recognize the existing RAID array at all. Therefore, the technician needs to reconfigure the RAID in order to restore operations with minimal downtime.

NEW QUESTION 14

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

```
ipconfig/all

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

Pinging 192.168.1.2 with 32 bytes of data:

Reply from 192.168.1.2: Request timed out
```

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

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

Answer: D

Explanation:

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

NEW QUESTION 15

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

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

Answer: D

Explanation:

The administrator should connect the server to a UPS to prevent this situation in the future. A UPS (Uninterruptible Power Supply) is a device that provides backup power to a server or other device in case of a power outage or surge. A UPS typically consists of one or more batteries and an inverter that converts the battery power into AC power that the server can use. A UPS can also protect the server from power fluctuations that can damage its components or cause data corruption. By connecting the server to a UPS, the administrator can ensure that the server will continue to run or shut down gracefully during a power failure.

NEW QUESTION 16

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

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

Answer: C

Explanation:

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

NEW QUESTION 20

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

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

Answer: A

Explanation:

To make a new drive available on a server, the administrator needs to partition the drive first. Partitioning is a process that divides the drive into one or more logical sections that can be formatted and assigned drive letters or mount points. Partitioning can be done using tools such as Disk Management on Windows or fdisk on Linux. Creating a new disk quota would not help, as disk quotas are used to limit the amount of disk space that users or groups can use on a partition. Configuring the drive as dynamic would not help either, as dynamic disks are used to create volumes that span multiple disks or use RAID features. Setting the

compression would not help, as compression is used to reduce the size of files on a partition. References:

<https://www.howtogeek.com/school/using-windows-admin-tools-like-a-pro/lesson2/><https://www.howtogeek.com/howto/17001/how-to-format-a-usb-drive-in-ubuntu-using-gparted/>

NEW QUESTION 23

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

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

Answer: C

Explanation:

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

NEW QUESTION 26

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

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

Answer: A

Explanation:

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

to Fibre Channel. References:

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

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

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

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

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

NEW QUESTION 28

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

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

Answer: D

Explanation:

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

NEW QUESTION 32

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 34

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 37

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 42

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 44

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

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

Answer: B

Explanation:

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

NEW QUESTION 46

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 49

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 52

An administrator discovers a Bash script file has the following permissions set in octal notation;
777

Which of the following is the MOST appropriate command to ensure only the root user can modify and execute the script?

- A. `chmod go-rw>`
- B. `chmod u=rwx`
- C. `chmod u+wx`
- D. `chmod g-rwx`

Answer: A

Explanation:

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

NEW QUESTION 53

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

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

Answer: A

Explanation:

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

NEW QUESTION 55

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 57

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

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

Answer: A

Explanation:

Two-person integrity is a security measure that can be implemented to increase security against a potential insider threat on an air-gapped system. An air-gapped system is a system that is isolated from any network connection and can only be accessed physically. An insider threat is a malicious actor who has authorized access to an organization's system or data and uses it for unauthorized or harmful purposes. Two-person integrity is a system of storage and handling that requires the presence of at least two authorized persons, each capable of detecting incorrect or unauthorized security procedures, for accessing certain sensitive data or material. This way, no single person can compromise the security or integrity of the data or material without being noticed by another person. SSO (Single Sign-On) is a feature that allows users to access multiple applications or systems with one set of credentials, but it does not prevent insider threats. SIEM (Security Information and Event Management) is a tool that collects and analyzes log data from various sources to detect and respond to security incidents, but it does not work on air-gapped systems. A Faraday cage is a structure that blocks electromagnetic signals from entering or leaving, but it does not prevent physical access or insider threats. MFA (Multi-Factor Authentication) is a method that requires users to provide two or more pieces of evidence to verify their identity, such as something they know, something they have, or something they are, but it does not prevent insider threats. References: <https://www.howtogeek.com/169080/air-gap-how-to-isolate-a-computer-to-protect-it-from-hackers/> <https://www.howtogeek.com/428483/what-is-end-to-end-encryption-and-why-does-it-matter/> <https://www.howtogeek.com/202794/what-is-the-difference-between-127.0.0.1-and-0.0.0.0/> <https://www.howtogeek.com/443611/how-to-encrypt-your-macs-system-drive-removable-devices-and-individual-files/>

NEW QUESTION 62

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 64

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

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

Answer: D

Explanation:

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

NEW QUESTION 68

An administrator is setting up a new server and has been asked to install an operating system that does not have a GUI because the server has limited resources. Which of the following installation options should the administrator use?

- A. Bare metal
- B. Headless
- C. Virtualized
- D. Slipstreamed

Answer: B

Explanation:

A headless installation is an installation method that does not require a graphical user interface (GUI) or a monitor, keyboard, and mouse. It can be done remotely through a network connection or a command-line interface. A headless installation is suitable for a server that has limited resources and does not need a GUI.

References:

? CompTIA Server+ Certification Exam Objectives1, page 14

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

NEW QUESTION 70

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

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

Answer: B

Explanation:

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

NEW QUESTION 75

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

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

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 79

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 80

Hosting data in different regional locations but not moving it for long periods of time describes:

- A. a cold site.
- B. data at rest.
- C. on-site retention.
- D. off-site storage.

Answer: B

Explanation:

Data at rest refers to data that is stored in a persistent state on any device or media, such as hard drives, tapes, or cloud storage. Data at rest does not move for long periods of time unless it is accessed or modified by authorized users or applications. A cold site (A) is a backup location that has minimal or no equipment and resources to resume business operations in case of a disaster. On-site retention © is a policy of keeping backup data on premises for a certain period of time before transferring it to an off-site location.

Off-site storage (D) is a method of storing backup data in a remote location that is physically or logically separated from the primary site. References:

<https://docs.microsoft.com/en-us/azure/security/fundamentals/encryption-atrest> <https://www.techopedia.com/definition/144/cold-site>

<https://www.enterprisestorageforum.com/backup/onsite-offsite-backup.html><https://www.techopedia.com/definition/24195/offsite-storage>

NEW QUESTION 83

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 84

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 memory¹. 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 85

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 89

A technician retails 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 90

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 91

A systems administrator has noticed performance degradation on a company file server, and one of the disks on it has a solid amber light. The administrator logs on to the disk utility and sees the array is rebuilding. Which of the following should the administrator do NEXT once the rebuild is finished?

- A. Restore the server from a snapshot.
- B. Restore the server from backup.
- C. Swap the drive and initialize the disk.
- D. Swap the drive and initialize the array.

Answer: C

Explanation:

The next action that the administrator should take once the rebuild is finished is to swap the drive and initialize the disk. This is to replace the faulty disk that has a solid amber light, which indicates a predictive failure or a SMART error. Initializing the disk will prepare it for use by the RAID controller and add it to the array. The administrator should also monitor the array status and performance after swapping the drive. Reference: <https://www.salvagedata.com/how-to-rebuild-a-failed-raid/>

NEW QUESTION 92

A server administrator is installing a new server on a manufacturing floor. Because the server is publicly accessible, security requires the server to undergo hardware hardening. Which of the following actions should the administrator take?

- A. Close unneeded ports.
- B. Disable unused services.
- C. Set a BIOS password.
- D. Apply driver updates.

Answer: C

Explanation:

An action that the administrator should take to harden the hardware of a new server is to set a BIOS password. BIOS (Basic Input/Output System) is a firmware that initializes the hardware components and settings of a system before loading the operating system. BIOS password is a security feature that requires a user to enter a password before accessing or modifying the BIOS settings or booting up the system. By setting a BIOS password, the administrator can prevent unauthorized or malicious users from changing the hardware configuration or boot order of the server.

References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 5, Lesson 5.1, Objective 5.1

NEW QUESTION 97

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 98

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 99

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 backup.

NEW QUESTION 103

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 107

A change in policy requires a complete backup of the accounting server every seven days and a backup of modified data every day. Which of the following would

be BEST to restore a full backup as quickly as possible in the event of a complete loss of server data?

- A. A full, weekly backup with daily open-file backups
- B. A full, weekly backup with daily archive backups
- C. A full, weekly backup with daily incremental backups
- D. A full, weekly backup with daily differential backups

Answer: D

Explanation:

A differential backup is a type of backup that copies all the files that have changed since the last full backup. A differential backup requires more storage space than an incremental backup, which only copies the files that have changed since the last backup of any type, but it also requires less time to restore in case of data loss. By combining a full, weekly backup with daily differential backups, the administrator can ensure that only two backup sets are needed to restore a full backup as quickly as possible. Verified References: [Incremental vs Differential Backup]

NEW QUESTION 109

A server administrator is configuring the IP address on a newly provisioned server in the testing environment. The network VLANs are configured as follows:

VLAN name	VLAN ID	Gateway IP address	Active switchports
Testing	10	192.168.10.1/24	2, 4, 6, 8, 10, 12, 14, 18
Production	20	192.168.20.1/24	3, 5, 7, 9, 11, 13, 15, 17
Administration	30	192.168.30.1/24	1, 24

The administrator configures the IP address for the new server as follows: IP address: 192.168.1.1/24

Default gateway: 192.168.10.1

A ping sent to the default gateway is not successful. Which of the following IP address/default gateway combinations should the administrator have used for the new server?

- A. IP address: 192.168.10.2/24 Default gateway: 192.168.10.1
- B. IP address: 192.168.1.2/24 Default gateway: 192.168.10.1
- C. IP address: 192.168.10.3/24 Default gateway: 192.168.20.1
- D. IP address: 192.168.10.24/24 Default gateway: 192.168.30.1

Answer: A

Explanation:

The IP address/default gateway combination that the administrator should have used for the new server is IP address: 192.168.10.2/24 and Default gateway: 192.168.10.1. The IP address and the default gateway of a device must be in the same subnet to communicate with each other. A subnet is a logical division of a network that allows devices to share a common prefix of their IP addresses. The subnet mask determines how many bits of the IP address are used for the network prefix and how many bits are used for the host identifier. A /24 subnet mask means that the first 24 bits of the IP address are used for the network prefix and the last 8 bits are used for the host identifier. Therefore, any IP address that has the same first 24 bits as the default gateway belongs to the same subnet. In this case, the default gateway has an IP address of 192.168.10.1/24, which means that any IP address that starts with 192.168.10.x/24 belongs to the same subnet. The new server has an IP address of 192.168.1.1/24, which does not match the first 24 bits of the default gateway, so it belongs to a different subnet and cannot communicate with the default gateway. To fix this issue, the administrator should change the IP address of the new server to an unused IP address that starts with 192.168.10.x/24, such as 192.168.10.2/24.

NEW QUESTION 112

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

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

Answer: D

Explanation:

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

NEW QUESTION 115

Which of the following BEST describes a guarantee of the amount of time it will take to restore a downed service?

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

Answer: A

Explanation:

RTO stands for Recovery Time Objective and it is a metric that defines the maximum acceptable amount of time that a system or service can be unavailable after a disaster or disruption. RTO is part of the business continuity planning and disaster recovery planning processes. RTO ensures a guarantee of the amount of time it will take to restore a downed service by setting a target or goal for recovery. RTO can vary depending on the criticality and priority of the service. References: <https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-objectives> (Objective 3.3)

NEW QUESTION 120

A company's servers are all displaying the wrong time. The server administrator confirms the time source is correct. Which of the following is MOST likely preventing the servers from obtaining the correct time?

- A. A firewall
- B. An antivirus
- C. AHIDS
- D. User account control

Answer: A

Explanation:

The most likely cause of the servers displaying the wrong time is A. A firewall. A firewall is a network security device that monitors and controls incoming and outgoing network traffic based on predefined rules. A firewall can block or allow certain ports, protocols, or applications that are used for network communication. One of the protocols that is used for time synchronization is the Network Time Protocol (NTP), which requires the use of UDP port 123 for all time synchronization¹. If a firewall blocks this port, it can prevent the servers from obtaining the correct time from the time source. Therefore, the server administrator should check the firewall settings and make sure that UDP port 123 is allowed for NTP traffic.

NEW QUESTION 121

A server in a remote datacenter is no longer responsive. Which of the following is the BEST solution to investigate this failure?

- A. Remote desktop
- B. Access via a crash cart
- C. Out-of-band management
- D. A Secure Shell connection

Answer: C

Explanation:

The best solution to investigate the failure of a server in a remote datacenter is out-of-band management. Out-of-band management is a method of accessing and controlling a server or a device using a dedicated channel that is separate from its normal network connection. Out-of-band management can use various technologies, such as serial ports, modems, KVM switches, or dedicated management cards or interfaces. Out-of-band management can provide remote access to servers or devices even when they are powered off, unresponsive, or disconnected from the network. Out-of-band management can enable troubleshooting, configuration, maintenance, or recovery tasks without requiring physical presence at the server location.

Reference:

https://www.lantronix.com/wp-content/uploads/pdf/Data_Center_Mgmt_WP.pdf

NEW QUESTION 126

Which of the following is the MOST secure method to access servers located in remote branch offices?

- A. Use an MFA out-of-band solution.
- B. Use a Telnet connection.
- C. Use a password complexity policy.
- D. Use a role-based access policy.

Answer: A

Explanation:

This is the most secure method to access servers located in remote branch offices because MFA stands for multi-factor authentication, which requires users to provide more than one piece of evidence to prove their identity. An out-of-band solution means that one of the factors is delivered through a separate channel, such as a phone call, a text message, or an email. This adds an extra layer of security and prevents unauthorized access even if a password is compromised. References: <https://docs.microsoft.com/en-us/azure/active-directory/authentication/concept-mfa-howitworks>

NEW QUESTION 129

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

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

Answer: A

Explanation:

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

NEW QUESTION 134

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

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

Answer: A

Explanation:

Bollards are an example of a method that is used to prevent motor vehicles from getting too close to building entrances and exits. Bollards are short, sturdy posts that are installed on sidewalks, parking lots, or roads to create physical barriers and control traffic flow. Bollards can be used to protect pedestrians, buildings, or other structures from vehicle collisions or attacks. Bollards can be made of various materials, such as metal, concrete, or plastic, and can be fixed, removable, or retractable.

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

NEW QUESTION 137

After the installation of an additional network card into a server, the server will not boot into the OS. A technician tests the network card in a different server with a different OS and verifies the card functions correctly. Which of the following should the technician do NEXT to troubleshoot this issue?

- A. Remove the original network card and attempt to boot using only the new network card.
- B. Check that the BIOS is configured to recognize the second network card.
- C. Ensure the server has enough RAM to run a second network card.
- D. Verify the network card is on the HCL for the OS.

Answer: D

Explanation:

The HCL stands for Hardware Compatibility List and it is a list of hardware devices that are tested and certified to work with a specific operating system. If a network card is not on the HCL for the OS, it may not function properly or cause compatibility issues. Therefore, verifying the network card is on the HCL for the OS should be the next step to troubleshoot this issue. References: <https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-objectives> (Objective 4.1)

NEW QUESTION 139

A backup application is copying only changed files each time it runs. During a restore, however, only a single file is used. Which of the following backup methods does this describe?

- A. Open file
- B. Synthetic full
- C. Full Incremental
- D. Full differential

Answer: B

Explanation:

A synthetic full backup is a backup method that describes copying only changed files each time it runs and using only a single file during a restore. A synthetic full backup is a backup approach that involves creating a new full backup by using the previous full backup and related incremental backups. This means that a backup solution does not have to transfer the full amount of data from the source machine and can synthesize the latest incremental backups with the last full backup to create a new full backup. This reduces the backup window and network bandwidth consumption. During a restore, only the latest synthetic full backup file is needed to recover the data. Open file backup is a backup method that allows backing up files that are in use or locked by applications. Full incremental backup is a backup method that involves performing a full backup first and then backing up only the changed files since the last backup. Full differential backup is a backup method that involves performing a full backup first and then backing up only the changed files since the last full backup. References: <https://www.nakivo.com/blog/what-is-synthetic-backup/> <https://www.howtogeek.com/192115/what-you-need-to-know-about-creating-system-image-backups/>

NEW QUESTION 141

Which of the following is a type of replication in which all files are replicated, all the time?

- A. Constant
- B. Application consistent
- C. Synthetic full
- D. Full

Answer: A

Explanation:

Constant replication is a type of replication in which all files are replicated, all the time. Replication is a process of copying data from one location to another for backup, recovery, or distribution purposes. Constant replication is also known as real-time replication or synchronous replication. It ensures that any changes made to the source data are immediately reflected on the target data without any delay or lag. Constant replication provides high availability and consistency, but it requires high bandwidth and low latency. Application consistent replication is a type of replication that ensures that the replicated data is consistent with the state of the application that uses it. It involves quiescing or pausing the application before taking a snapshot of the data and resuming the application after the snapshot is taken. Application consistent replication provides better recovery point objectives than crash consistent replication, which does not quiesce the application before taking a snapshot. Synthetic full replication is a type of replication that involves creating a new full backup by using the previous full backup and related incremental backups. It reduces the backup window and network bandwidth consumption by transferring only changed data from the source to the target. Full replication is a type of replication that involves copying all data from the source to the target regardless of whether it has changed or not. It provides a complete backup of the data, but it requires more storage space and network bandwidth than incremental or differential replication. References: <https://www.howtogeek.com/199068/how-to-upgrade-your-existing-hard-drive-in-under-an-hour/> <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 145

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

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

Answer: B

Explanation:

The cluster feature that will allow Ann to complete her task is failback. A cluster is a group of servers that work together to provide high availability, scalability, and load balancing for applications or services. A cluster can have different nodes or members that have different roles or states. An active node is a node that is currently running an application or service and serving requests from clients. A passive node is a node that is on standby and ready to take over if the active node fails. A failover is a process of switching from a failed or unavailable node to another node in a cluster. A failback is a process of switching back from a failover node to the original node after it becomes available again. Failback can be automatic or manual depending on the cluster configuration.

NEW QUESTION 148

A technician learns users are unable to log in to a Linux server with known-working LDAP credentials. The technician logs in to the server with a local account and confirms the system is functional can communicate over the network, and is configured correctly. However, the server log has entries regarding Kerberos errors. Which of the following is the MOST likely source of the issue?

- A. A local firewall is blocking authentication requests.
- B. The users have expired passwords
- C. The system clock is off by more than five minutes
- D. The server has no access to the LDAP host

Answer: C

Explanation:

Kerberos is a network authentication protocol that uses tickets to allow clients and servers to prove their identity to each other. Kerberos relies on accurate time synchronization between the parties involved, as the tickets have expiration dates and timestamps. If the system clock of a Linux server is off by more than five minutes from the LDAP server or the domain controller, the Kerberos authentication will fail and generate errors. A local firewall is unlikely to block authentication requests if the server can communicate over the network and is configured correctly. The users' passwords are not relevant if they are known-working LDAP credentials. The server has access to the LDAP host if it can communicate over the network and is configured correctly. References:

- ? https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/identity_management_guide/kerberos_errors
- ? <https://www.ibm.com/docs/en/aix/7.2?topic=authentication-kerberos-time-synchronization>

NEW QUESTION 153

Which of the following describes a configuration in which both nodes of a redundant system respond to service requests whenever possible?

- A. Active-passive
- B. Failover
- C. Active-active
- D. Fallback

Answer: C

Explanation:

Active-active is a configuration in which both nodes of a redundant system respond to service requests whenever possible. It can improve the performance, availability, and load balancing of the system by distributing the workload among the nodes. However, it also requires more synchronization and coordination between the nodes to avoid conflicts or errors. Verified References: [Active-active], [Redundant system]

NEW QUESTION 154

Which of the following would be BEST to help protect an organization against social engineering?

- A. More complex passwords
- B. Recurring training and support
- C. Single sign-on
- D. An updated code of conduct to enforce social media

Answer: B

Explanation:

The best way to protect an organization against social engineering is to provide recurring training and support. Social engineering is a type of attack that exploits human psychology and behavior to manipulate people into divulging confidential information or performing malicious actions. Social engineering can take various forms, such as phishing emails, phone calls, impersonation, baiting, or quid pro quo. The best defense against social engineering is to educate and empower the employees to recognize and avoid common social engineering techniques and report any suspicious activities or incidents. Recurring training and support can help raise awareness and reinforce best practices among the employees.

NEW QUESTION 156

A company has a data center that is located at its headquarters, and it has a warm site that is located 20mi (32km) away, which serves as a DR location. Which of the following should the company design and implement to ensure its DR site is adequate?

- A. Set up the warm site as a DR cold site.
- B. Set up a DR site that is in the cloud and in the same region.
- C. Set up the warm site as a DR hot site.
- D. Set up a DR site that is geographically located in another region.

Answer: D

Explanation:

A DR site is a backup site that can be used to restore business operations in case of a disaster that affects the primary site. A warm site is a DR site that has some equipment and data ready to be activated quickly, but not as fast as a hot site that has fully operational systems and data. A cold site is a DR site that has only basic infrastructure and no equipment or data. The location of a DR site is an important factor to consider when designing and implementing a DR plan. A DR site that is too close to the primary site may be affected by the same disaster, such as a power outage, a flood, or an earthquake. A DR site that is too far away

from the primary site may incur higher costs and latency issues. Therefore, a good practice is to set up a DR site that is geographically located in another region that has different risk factors and environmental conditions than the primary site. This can help ensure that the DR site is available and accessible when needed. References: <https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-objectives> (Objective 3.3)

NEW QUESTION 157

Users at a company are licensed to use an application that is restricted by the number of active sessions. Which of the following best describes this licensing model?

- A. Per-server
- B. per-seat
- C. Per-concurrent user
- D. per-core

Answer: C

Explanation:

The per-concurrent user licensing model is a type of licensing model that restricts the number of active sessions or connections to a software application at any given time. This means that multiple users can share the same license, as long as they do not access the application simultaneously. This model is often used for applications that are accessed intermittently or for a short duration by different users, such as remote access software, web-based applications, or testing tools.

NEW QUESTION 158

Which of the following script types uses commands that start with sec-?

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

Answer: C

Explanation:

PowerShell is a scripting language and a command-line shell that uses commands that start with sec- to perform security-related tasks. For example, sec-edit is a command that edits security policies, sec-logon is a command that manages logon sessions, and sec-policy is a command that applies security templates. Verified References: [PowerShell security commands], [Security policy]

NEW QUESTION 162

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

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

Answer: D

Explanation:

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

NEW QUESTION 163

Which of the following should be configured in pairs on a server to provide network redundancy?

- A. MRU
- B. SCP
- C. DLP
- D. CPU
- E. NIC

Answer: E

Explanation:

NIC stands for network interface card, which is a hardware component that allows a server to connect to a network. Configuring NICs in pairs on a server would provide network redundancy, meaning that if one NIC fails, the other one can take over and maintain network connectivity. The other options are not related to network redundancy.

NEW QUESTION 166

Which of the following backup types only records changes to the data blocks on a virtual machine?

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

Answer: B

Explanation:

The backup type that only records changes to the data blocks on a virtual machine is snapshot. 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 167

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

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

Answer: D

Explanation:

The concept of right to downgrade allows a previous version of an OS to be installed and covered by the same license as the newer version. For example, if a customer has a license for Windows 10 Pro, they can choose to install Windows 8.1 Pro or Windows 7 Professional instead and still be compliant with the license terms. Downgrade rights are granted by Microsoft for certain products and programs, such as Windows and Windows Server software acquired through Commercial Licensing, OEM, or retail channels. Downgrade rights are intended to provide customers with flexibility and compatibility when using Microsoft software.

NEW QUESTION 172

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

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

Answer: A

Explanation:

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

NEW QUESTION 173

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

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

Answer: D

Explanation:

The administrator should rebuild the volume to restore it after replacing the failed hard drive. A volume is a logical unit of storage that can span across multiple physical disks. A volume can be configured with different levels of RAID (Redundant Array of Independent Disks) to provide fault tolerance and performance enhancement. When a hard drive in a RAID volume fails, the data on that drive can be reconstructed from the remaining drives using parity or mirroring techniques. However, this process requires a new hard drive to replace the failed one and a rebuild operation to copy the data from the existing drives to the new one. Rebuilding a volume can take a long time depending on the size and speed of the drives and the RAID level.

NEW QUESTION 174

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

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

Answer: B

Explanation:

A subnet mask is a number that distinguishes the network address and the host address within an IP address¹. A subnet mask allows network traffic to understand IP addresses by splitting them into the network and host addresses. If the subnet mask is incorrect, the network traffic may not be able to determine the correct destination for the packets, and only reach some of the servers that have similar IP addresses. For example, if the new virtual server has an IP address of 192.168.1.100 and a subnet mask of 255.255.0.0, it can only communicate with servers that have IP addresses in the range of 192.168.0.0 to 192.168.255.255. To fix this issue, the systems administrator needs to check and correct the subnet mask of the new virtual server according to the network configuration.

NEW QUESTION 175

An administrator is configuring a host-based firewall for a server. The server needs to allow SSH, FTP, and LDAP traffic. Which of the following ports must be configured so this traffic will be allowed? (Select THREE).

- A. 21
- B. 22
- C. 53
- D. 67
- E. 69
- F. 110
- G. 123
- H. 389

Answer: ABH

Explanation:

These are the port numbers that must be configured on a host-based firewall for a server that needs to allow SSH, FTP, and LDAP traffic. A port number is a numerical identifier that specifies a communication endpoint for a network protocol or an application. A host-based firewall is a software tool that monitors and controls incoming and outgoing network traffic on a single host based on predefined rules. SSH (Secure Shell) is a protocol that allows secure remote access and file transfer over an encrypted connection. The default port number for SSH is 22. FTP (File Transfer Protocol) is a protocol that allows transferring files between hosts over a network connection. The default port number for FTP is 21. LDAP (Lightweight Directory Access Protocol) is a protocol that allows accessing and managing directory services over a network connection. The default port number for LDAP is 389. References: <https://www.howtogeek.com/190014/virtualization-basics-understanding-techniques-and-fundamentals/> <https://www.howtogeek.com/220152/what-is-the-difference-between-127.0.0.1-and-0.0.0.0/><https://www.howtogeek.com/428483/what-is-end-to-end-encryption-and-why-does-it-matter/>

NEW QUESTION 178

The network's IDS is giving multiple alerts that unauthorized traffic from a critical application server is being sent to a known-bad public IP address.

One of the alerts contains the following information: Exploit Alert

Attempted User Privilege Gain 2/2/07-3: 09:09 10.1.200.32

--> 208.206.12.9:80

This server application is part of a cluster in which two other servers are also servicing clients. The server administrator has verified the other servers are not sending out traffic to that public IP address. The IP address subnet of the application servers is 10.1.200.0/26. Which of the following should the administrator perform to ensure only authorized traffic is being sent from the application server and downtime is minimized? (Select two).

- A. Disable all services on the affected application server.
- B. Perform a vulnerability scan on all the servers within the cluster and patch accordingly.
- C. Block access to 208.206.12.9 from all servers on the network.
- D. Change the IP address of all the servers in the cluster to the 208.206.12.0/26 subnet.
- E. Enable GPO to install an antivirus on all the servers and perform a weekly reboot.
- F. Perform an antivirus scan on all servers within the cluster and reboot each server.

Answer: BF

Explanation:

The administrator should perform an antivirus scan on all servers within the cluster and reboot each server, and block access to 208.206.12.9 from all servers on the network. These actions will help to remove any malware that may have infected the application server and prevent any further unauthorized traffic to the known-bad public IP address. An antivirus scan can detect and remove malicious software that may be sending data to an external source, and a reboot can clear any temporary files or processes that may be related to the malware. Blocking access to 208.206.12.9 from all servers on the network can prevent any future attempts to communicate with the malicious IP address.

References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 3, Lesson 3.4, Objective 3.4; Chapter 6, Lesson 6.2, Objective 6.2

NEW QUESTION 182

An administrator is configuring a new server for use as a database server. It will have two mirrored drives to hold the operating system, and there will be three drive bays remaining for storage. Which of the following RAID levels will yield the BEST combination of available space and redundancy?

- A. RAID
- B. RAID 1
- C. RAIDS
- D. RAID 10

Answer: D

Explanation:

RAID 10 is the RAID level that will yield the best combination of available space and redundancy when configuring a new server for use as a database server with two mirrored drives for the operating system and three drive bays remaining for storage. RAID 10, also known as RAID 1+0, is a RAID configuration that combines disk mirroring and disk striping to protect data. It requires a minimum of four disks and stripes data across mirrored pairs. As long as one disk in each mirrored pair is functional, data can be retrieved. RAID 10 provides high performance, fault tolerance, and fast recovery, but it reduces storage capacity by half. RAID 0 is a RAID configuration that splits data across two or more drives without parity or redundancy. It improves performance but offers no fault tolerance. If one drive fails in RAID 0, all data is lost and the system cannot boot. RAID 1 is a RAID configuration that duplicates data across two or more drives. It provides fault tolerance and improves read performance, but reduces storage capacity by half. If one drive fails in RAID 1, the other drive can continue to operate without data loss or system downtime. RAID 5 is a RAID configuration that stripes data across three or more drives with parity information. It provides fault tolerance and improves performance, but reduces storage capacity by one drive's worth of space. RAID 5 can tolerate one drive failure without data loss, but not two or more. References: <https://www.howtogeek.com/199068/how-to-upgrade-your-existing-hard-drive-in-under-an-hour/><https://www.howtogeek.com/202794/what-is-the-difference-between-127.0.0.1-and-0.0.0.0/>

NEW QUESTION 186

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

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

Answer: B

Explanation:

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

Reference:

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

NEW QUESTION 189

A server administrator is gathering business requirements to determine how frequently backups need to be performed on an application server. Which of the following is the administrator attempting to establish?

- A. MTBF
- B. RPO
- C. MTTR
- D. RFC

Answer: B

Explanation:

The administrator is attempting to establish the recovery point objective (RPO) by determining how frequently backups need to be performed on an application server. RPO is a metric that defines how much data can be lost or how far back in time a recovery can go in case of a disaster or disruption, based on the business requirements and impact analysis of an organization or system. RPO is measured by the time interval between backups or snapshots of data, such as hourly, daily, weekly, etc., depending on how critical or sensitive the data is and how often it changes or updates. References: [CompTIA Server+ Certification Exam Objectives], Domain 5.0: Security, Objective 5.3: Given a scenario, explain methods and techniques to secure data.

NEW QUESTION 193

A technician needs to restore data from a backup. The technician has these files in the backup inventory:

Name	Size
01012020.bak	100MB
01022020.bak	10MB
01032020.bak	5MB
01042020.bak	7MB
01052020.bak	120MB
01062020.bak	8MB
01072020.bak	10MB

Which of the following backup types is being used if the file 01062020.bak requires another file to restore data?

- A. Full
- B. Incremental
- C. Snapshot
- D. Differential

Answer: B

Explanation:

An incremental backup only backs up files that have changed since the last backup, whether it was a full or an incremental backup. Therefore, an incremental backup file may require another file to restore data, depending on the sequence of backups. A full backup backs up all files and does not require any other file to restore data. A snapshot is a point-in-time copy of data that does not depend on other files. A differential backup backs up files that have changed since the last full backup and does not require any other file to restore data.

NEW QUESTION 196

Which of the following is a system that scans outgoing email for account numbers, sensitive phrases, and other forms of PII?

- A. SIEM
- B. DLP
- C. HIDS
- D. IPS

Answer: B

Explanation:

DLP stands for Data Loss Prevention and it is a system that scans outgoing email for account numbers, sensitive phrases, and other forms of PII (Personally Identifiable Information). DLP can help prevent data breaches, comply with regulations, and protect the privacy of customers and employees. DLP can also block, encrypt, or quarantine emails that contain sensitive data. References: <https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-objectives> (Objective 3.2)

NEW QUESTION 201

Which of the following can be used to map a network drive to a user profile?

- A. System service

- B. Network service
- C. Login script
- D. Kickstart script

Answer: C

Explanation:

A login script is a file that contains commands or instructions that are executed when a user logs into a system or network. A login script can be used to map a network drive to a user profile, which means that the user will have access to a shared folder or resource on another computer or server. A login script can be written in various languages, such as batch, PowerShell, or VBScript, and can be assigned to a user or a group using tools such as Group Policy or Active Directory. A system service (A) is a program that runs in the background and performs tasks that are essential for the operation of the system, such as security, networking, or hardware management. A system service does not map a network drive to a user profile. A network service (B) is a program that provides functionality or resources to other programs or devices over a network, such as file sharing, printing, or web hosting. A network service does not map a network drive to a user profile. A kickstart script (D) is a file that contains configuration settings and commands for automated installation of Linux operating systems. A kickstart script does not map a network drive to a user profile. References: <https://www.howtogeek.com/118452/how-to-map-network-drives-from-the-command-prompt-in-windows/> <https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/logon>

NEW QUESTION 204

Which of the following ensures a secondary network path is available if the primary connection fails?

- A. Link aggregation
- B. Most recently used
- C. Heartbeat
- D. Fault tolerance

Answer: D

Explanation:

Fault tolerance is the ability of a system to continue functioning in the event of a failure of one or more of its components. Fault tolerance can ensure a secondary network path is available if the primary connection fails. Fault tolerance can be achieved by using redundant components, such as network cards, cables, switches, routers, etc., that can take over the function of the failed component without interrupting the service. References: <https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-objectives> (Objective 2.2)

NEW QUESTION 207

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual SK0-005 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the SK0-005 Product From:

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

Money Back Guarantee

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