ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

Which of the following is typical of software licensing in the cloud?

A. Per socket
B. Perpetual
C. Subscription-based
D. Site-based

Answer: C

Explanation:

Cloud software licensing refers to the process of managing and storing software licenses in the cloud. The benefits of cloud software licensing models are vast. The main and most attractive benefit has to do with the ease of use for software vendors and the ability to provide customizable cloud software license management based on customer needs and desires1. Cloud-based licensing gives software developers and vendors the opportunity to deliver software easily and quickly and gives customers full control over their licenses, their analytics, and more1. Cloud based licensing gives software sellers the ability to add subscription models to their roster of services1. Subscription models are one of the most popular forms of licensing today1. Users sign up for a subscription (often based on various options and levels of use, features, etc.) and receive theirlicenses instantly1.
Reference: 1 Everything You Need to Know about Cloud Licensing | Thales

Question: 2

A server administrator wants to run a performance monitor for optimal system utilization. Which of the following metrics can the administrator use for monitoring? (Choose two.)

A. Memory
B. Page file
C. Services
D. Application
E. CPU
F. Heartbeat

Answer: A,E

Explanation:

Memory and CPU are two metrics that can be used for monitoring system utilization. Memory refers to the amount of RAM that is available and used by the system and its processes. CPU refers to the percentage of processor time that is consumed by the system and its processes. Both memory and CPU can affect the performance and responsiveness of the system and its applications. Monitoring memory and CPU can help identify bottlenecks, resource contention, memory leaks, high load, etc.

Question: 3

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

```
ipconfig/all

IPv4 address: 192.168,1.5
Subnet mask: 255.255.255.0
Default gateway: 192.168.1.2

Pinging 192.168.1.2 with 32 bytes of data

   Reply from 192.168.1.2: Request timed out
   Reply from 192.168.1.2: Request timed out
   Reply from 192.168.1.2: Request timed out
   Reply from 192.168.1,2: Request timed out
```

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

A. Network port security

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

Answer: D

Explanation:

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

the NIC and the switch.

Question: 4

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

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

Answer: C

Explanation:

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

Question: 5

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

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

Answer: B,F

Explanation:

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

Question: 6

A company has implemented a requirement to encrypt all the hard drives on its servers as part of a data loss prevention strategy. Which of the following should the company also perform as a data loss prevention method?

A. Encrypt all network traffic
B. Implement MFA on all the servers with encrypted data
C. Block the servers from using an encrypted USB
D. Implement port security on the switches

Answer: B

Explanation:

The company should also implement MFA on all the servers with encrypted data as a data loss prevention method. MFA stands for multi-factor authentication, which is a method of verifying a user's identity by requiring two or more pieces of evidence, such as something they know (e.g., a password), something they have (e.g., a token), or something they are (e.g., a fingerprint). MFA adds an extra layer of security to prevent unauthorized access to sensitive data, even if the user's password is compromised or stolen. Encrypting the hard drives on the servers protects the data from being read or copied if the drives are physically removed or stolen, but it does not prevent unauthorized access to the data if the user's credentials are valid.

Question: 7

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

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

Answer: C

Explanation:

The administrator should use 172.16.19.241 as an IP address to be in compliance with RFC 1918 standard. RFC 1918 defines three ranges of IP addresses that are reserved for private internets, meaning they are not globally routable on the public Internet and can be used within an enterprise without any risk of conflict or overlap with other networks. These ranges are:

10.0.0.0 - 10.255.255.255 (10/8 prefix) 172.16.0.0 - 172.31.255.255 (172.16/12 prefix) 192.168.0.0 - 192.168.255.255 (192.168/16 prefix)

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

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

Question: 8

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.

Question: 9

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 itsconfiguration, 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.

Question: 10

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

dr-xr-xr-- /home/Ann

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

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

Answer: D

Explanation:

The administrator should use the command chmod 754 /home/Ann to resolve the issue without granting unnecessary permissions. The chmod command is used to change the permissions of files and directories on a Linux server. The

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

Reference:

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

Question: 11

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.

Question: 12

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

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

Answer: D

Explanation:

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

Question: 13

A systems administrator is preparing to install two servers in a single rack. The administrator is concerned that having both servers in one rack will increase the chance of power issues due to the increased load. Which of the following should the administrator implement FIRST to address the issue?

A. Separate circuits
B. An uninterruptible power supply
C. Increased PDU capacity
D. Redundant power supplies

Answer: A

Explanation:

The administrator should implement separate circuits first to address the issue of power issues due to the increased load. Separate circuits are electrical wiring systems that provide independent power sources for different devices or groups of devices. By using separate circuits, the administrator can avoid overloading a single circuit with too many servers and reduce the risk of power outages, surges, or fires. Separate circuits also provide redundancy and fault tolerance, as a failure in one circuit will not affect the other circuit.

Question: 14

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

Question: 15

A technician is installing a variety of servers in a rack. Which of the following is the BEST course of action for the technician to take while loading the rack?

A. Alternate the direction of the airflow
B. Install the heaviest server at the bottom of the rack
C. Place a UPS at the top of the rack
D. Leave 1U of space between each server

Answer: B

Explanation:

The technician should install the heaviest server at the bottom of the rack to load the rack properly. Installing the heaviest server at the bottom of the rack helps to balance the weight distribution and prevent the rack from tipping over or collapsing. Installing the heaviest server at the bottom of the rack also makes it easier to access and service the server without lifting or moving it. Installing the heaviest server at any other position in the rack could create instability and safety hazards.

Question: 16

A technician is configuring a server that requires secure remote access. Which of the following ports should the technician use?

A. 21
B. 22
C. 23
D. 443

Answer: B

Explanation:

The technician should use port 22 to configure a server that requires secure remote access. Port 22 is the default port for Secure Shell (SSH), which is a protocol that allows secure remote login and command execution over a network connection using a command-line interface (CLI). SSH encrypts both the authentication and data transmission between the client and the server, preventing eavesdropping, tampering, or spoofing. SSH can be used to perform various tasks on a server remotely, such as configuration, administration, maintenance, troubleshooting, etc.

Question: 17

A server administrator is using remote access to update a server. The administrator notices numerous error messages when using YUM to update the applications on a server. Which of the following should the administrator check FIRST?

A. Network connectivity on the server
B. LVM status on the server

C. Disk space in the /var directory
D. YUM dependencies

Answer: C

Explanation:

The administrator should check disk space in the /var directory first when using YUM to update applications on a server. YUM stands for Yellowdog Updater Modified, which is a software package manager for Linux systems that use RPM (Red HatPackage Manager) packages. YUM downloads and installs packages from online repositories and resolves dependencies automatically. YUM stores its cache files in the /var/cache/yum directory by default. These cache files include metadata and package data for each repository that YUM uses. If there is not enough disk space in the /var directory, YUM may fail to update applications and generate error messages.

Question: 18

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

Question: 19

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

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

Answer: B

Explanation:

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

Question: 20

Which of the following tools will analyze network logs in real time to report on suspicious log events?

A. Syslog
B. DLP
C. SIEM
D. HIPS

Answer: C

Explanation:

SIEM is the tool that will analyze network logs in real time to report on suspicious log events. SIEM stands for Security Information and Event Management, which is a software solution that collects, analyzes, and correlates log data from various sources, such as servers, firewalls, routers, antivirus software, etc. SIEM can detect anomalies, patterns, trends, and threats in the log data and generate alerts or reports for security monitoring and incident response. SIEM can also provide historical analysis and compliance reporting for audit purposes.

Reference:

https://www.manageengine.com/products/eventlog/syslog-server.html

Question: 21

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

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

Answer: B

Explanation:

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

name is John, the script will be mapped to \server\John.

Reference:

https://social.technet.microsoft.com/Forums/windows/en-US/07cfcb73-796d-48aa-96a9-08280a1ef25a/ mapping-home-directory-with-username-variable?forum=w7itprogeneral

Question: 22

A server that recently received hardware upgrades has begun to experience random BSOD conditions. Which of the following are likely causes of the issue? (Choose two.)

A. Faulty memory
B. Data partition error
C. Incorrectly seated memory
D. Incompatible disk speed
E. Uninitialized disk
F. Overallocated memory

Answer: A,C

Explanation:

Faulty memory and incorrectly seated memory are likely causes of the random BSOD conditions on the server. Memory is one of the most common hardware components that can cause BSOD (Blue Screen of Death) errors on Windows systems. BSOD errors occur when the system encounters a fatal error that prevents it from continuing to operate normally. Memory errors can be caused by faulty or incompatible memory modules that have physical defects or manufacturing flaws. Memory errors can also be caused by incorrectly seated memory modules that are not properly inserted or locked into the memory slots on the motherboard. This can result in loose or poor connections between the memory modules and the motherboard.

Question: 23

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

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

Answer: B

Explanation:

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

Question: 24

A technician is attempting to update a server's firmware. After inserting the media for the firmware and restarting the server, the machine starts normally into the OS. Which of the following should the technician do NEXT to install the firmware?

A. Press F8 to enter safe mode
B. Boot from the media
C. Enable HIDS on the server
D. Log in with an administrative account

Answer: B

Explanation:

The technician should boot from the media to install the firmware on the server. Firmware is a type of software that controls the low-level functions of hardware devices, such as BIOS (Basic Input/Output System), RAID controllers, network cards, etc. Firmware updates are often provided by hardware manufacturers to fix bugs, improve performance, or add new features to their devices. To install firmware updates on a server, the technician needs to boot from a media device (such as a CD- ROM, DVD-ROM, USB flash drive, etc.) that contains the firmware files and installationprogram. The technician cannot install firmware updates from within the operating system because firmware updates often require restarting or

resetting the hardware devices.

Question: 25

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 readwrite 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/QUESTIO N NO :s/149399/how-to-remount-as-read-write-a-specific- mount-of-device

Question: 26

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.

Question: 27

A server administrator needs to harden a server by only allowing secure traffic and DNS inquiries. A port scan reports the following ports are open:

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

Answer: D

Explanation:

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

Question: 28

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

A. 21
B. 22
C. 23

D. 53
E. 443
F. 636

Answer: A,C

Explanation:

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

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

Question: 29

Which of the following must a server administrator do to ensure data on the SAN is not compromised if it is leaked?

A. Encrypt the data that is leaving the SAN
B. Encrypt the data at rest
C. Encrypt the host servers
D. Encrypt all the network traffic

Answer: B

Explanation:

The administrator must encrypt the data at rest to ensure data on the SAN is not compromised if it is leaked. Data at rest refers to data that is stored on a device or a medium, such as a hard drive, a flash

drive, or a SAN (Storage Area Network). Data at rest can be leaked if the device or the medium is lost, stolen, or accessed byunauthorized parties. Encrypting data at rest means applying an algorithm that transforms the data into an unreadable format that can only be decrypted with a key. Encryption protects data at rest from being exposed or misused by attackers who may obtain the device or the medium.

Question: 30

A server technician has been asked to upload a few files from the internal web server to the internal FTP server. The technician logs in to the web server using PuTTY, but the connection to the FTP server fails. However, the FTP connection from the technician's workstation is successful. To troubleshoot the issue, the technician executes the following command on both the web server and the workstation:

ping ftp.acme.local

The IP address in the command output is different on each machine. Which of the following is the MOST likely reason for the connection failure?

A. A misconfigured firewall
B. A misconfigured hosts.deny file
C. A misconfigured hosts file
D. A misconfigured hosts.allow file

Answer: D

Explanation:

A misconfigured hosts file can cause name resolution issues on a server. A hosts file is a text file that maps hostnames to IP addresses on a local system. It can be used to override DNS settings or provide custom name resolution for testing purposes. However, if the hosts file contains incorrect or outdated entries, it can prevent the system from resolving hostnames properly and cause connectivity problems. To fix this issue, the administrator should check and edit the hosts file accordingly.

Question: 31

A company deploys antivirus, anti-malware, and firewalls that can be assumed to be functioning properly. Which of the following is the MOST likely system vulnerability?

A. Insider threat
B. Worms
C. Ransomware
D. Open ports
E. Two-person integrity

Answer: A

Explanation:

Insider threat is the most likely system vulnerability in a company that deploys antivirus, antimalware, and firewalls that can be assumed to be functioning properly. An insider threat is a malicious or negligent act by an authorized user of a system or network that compromises the security or integrity of the system or network. An insider threat can include data theft, sabotage, espionage, fraud, or other types of attacks. Antivirus, anti-malware, and firewalls are security tools that can protect a system or network from external threats, such as viruses, worms, ransomware, or open ports. However, these tools cannot prevent an insider threat from exploiting their access privileges or credentials to harm the system or network.

Question: 32

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

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

Answer: D

Explanation:

A port scanner is the tool that would most likely help the analyst determine if the applications are running on a remote server. A port scanner is a software tool that scans a network device for open ports. Ports are logical endpoints for network communication that are associated with specific

applications or services. By scanning the ports on a remote server, the analyst can identify what applications or services are running on that server and what protocols they are using. A port scanner can also help detect potential vulnerabilities or misconfigurations on a server.

Question: 33

A server is performing slowly, and users are reporting issues connecting to the application on that server. Upon investigation, the server administrator notices several unauthorizedservices running on that server that are successfully communicating to an external site. Which of the following are MOST likely causing the issue?
(Choose two.)

A. Adware is installed on the users' devices

B. The firewall rule for the server is misconfigured
C. The server is infected with a virus
D. Intrusion detection is enabled on the network
E. Unnecessary services are disabled on the server
F. SELinux is enabled on the server

Answer: C,F

Explanation:

The server is infected with a virus and SELinux is enabled on the server are most likely causing the issue of unauthorized services running on the server. A virus is a type of malicious software that infects a system and performs unwanted or harmful actions, such as creating, modifying, deleting, or executing files. A virus can also create backdoors or open ports on a system to allow remote access or communication with external sites. SELinux (Security-Enhanced Linux) is a security module for Linux systems that enforces mandatory access control policies on processes and files. SELinux can prevent unauthorized services from running on a server by restricting their access to resources based on their security context. However, SELinux can also cause problems if it is not configured properly or if it conflicts with other security tools.

Question: 34

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 areused 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.

Question: 35

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.

Question: 36

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

unauthorized server access? (Choose two.)

A. Audit all group privileges and permissions
B. Run a checksum tool against all the files on the server
C. Stop all unneeded services and block the ports on the firewall

D. Initialize a port scan on the server to identify open ports
E. Enable port forwarding on port 80
F. Install a NIDS on the server to prevent network intrusions

Answer: C,F

Explanation:

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

Question: 37

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.

Question: 38

A technician runs top on a dual-core server and notes the following conditions: top –- 14:32:27, 364 days, 14 usersload average 60.5 12.4 13.6

Which of the following actions should the administrator take?

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

Answer: C

Explanation:

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

Question: 39

A technician needs to set up a server backup method for some systems. The company's management team wants to have quick restores but minimize the amount of backup media required. Which of the following are

the BEST backup methods to use to support the management's priorities? (Choose two.)

A. Differential
B. Synthetic full
C. Archive
D. Full
E. Incremental
F. Open file

Answer: A,E

Explanation:

The best backup methods to use to support the management's priorities are differential and incremental. A backup is a process of copying data from a source to a destination for the purpose of restoring it in case of data loss or corruption. There are different types of backup methods that vary in terms of speed, efficiency, and storage requirements. Differential and incremental backups are two types of partial backups that only copy the data that has changed since the last full backup. A full backup is a type of backup that copies all the data from the source to the destination. A full backup provides the most complete and reliable restore option, but it also takes the longest time and requires the most storage space. A

differential backup copies only the data that has changed since the last full backup. A differential backup provides a faster restore option than an incremental backup, but it also takes more time and requires more storage space than an incremental backup. An incremental backup copies only the data that has changed since the last backup, whether it was a full or an incremental backup. An incremental backup provides the fastest and most efficient backup option, but it also requires multiple backups to restore the data completely.

Question: 40

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. Afailback 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.

Question: 41

Which of the following policies would be BEST to deter a brute-force login attack?

A. Password complexity
B. Password reuse
C. Account age threshold
D. Account lockout threshold

Answer: D

Explanation:

The best policy to deter a brute-force login attack is account lockout threshold. A brute-force login attack is a type of attack that tries to guess a user's password by trying different combinations of characters until it finds the correct one. This attack can be performed manually or with automated tools that use dictionaries, wordlists, or algorithms. An account lockout threshold is a policy that specifies how many failed login attempts are allowed before an account is locked out temporarily or permanently. This policy prevents an attacker from trying unlimited password guesses and reduces the chances of finding the correct password.

Question: 42

A technician is trying to determine the reason why a Linux server is not communicating on a network. The returned network configuration is as follows:

eth0: flags=4163<UP, BROADCAST,RUNNING,MULTICAST> mtu 1500 inet 127.0.0.1 network 255.255.0.0 broadcast 127.0.0.1

Which of the following BEST describes what is happening?

A. The server is configured to use DHCP on a network that has multiple scope options
B. The server is configured to use DHCP, but the DHCP server is sending an incorrect subnet mask
C. The server is configured to use DHCP on a network that does not have a DHCP server
D. The server is configured to use DHCP, but the DHCP server is sending an incorrect MTU setting

Answer: C

Explanation:

The reason why the Linux server is not communicating on a network is that it is configured to use DHCP on a network that does not have a DHCP server. DHCP (Dynamic Host Configuration Protocol) is a protocol that allows a client device to obtain an IP address and other network configuration parameters from a DHCP server automatically. However, if there is no DHCP server on the network, the client device will not be able to obtain a valid IP address and will assign itself a link-local address instead. A link-local address is an IP address that is only valid within a local network segment and cannot be used for communication outside of it. A link-local address has a prefix of 169.254/16 in IPv4 or fe80::/10 in IPv6. In this case, the Linux server has assigned itself a link-local address of 127.0.0.1, which is also known as the loopback address. The loopback address is used for testing and troubleshooting purposes and refers to the device itself. It cannot be used for communication with other devices on the network.

Question: 43

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

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

Answer: C

Explanation:

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

Reference:

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

Question: 44

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

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

Answer: D

Explanation:

The administrator should use SCP to transfer log files from a Linux server to a Windows workstation. SCP (Secure Copy

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

Question: 45

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.

Question: 46

Which of the following describes the installation of an OS contained entirely within another OS installation?

A. Host
B. Bridge
C. Hypervisor
D. Guest

Answer: D

Explanation:

The installation of an OS contained entirely within another OS installation is described as a guest. A guest is a term that refers to a virtual machine (VM) that runs on top of a host operating system (OS)

using a hypervisor or a virtualization software. A guest can have a different OS than the host, and can run multiple applications or services independently from the host. A guest can also be isolated from the host and other guests for security or testing purposes.

Question: 47

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

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

Answer: B

Explanation:

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

Question: 48

An administrator is configuring a server that will host a high-performance financial application.
Which of the following disk types will serve this purpose?

A. SAS SSD
B. SATA SSD
C. SAS drive with 10000rpm
D. SATA drive with 15000rpm

Answer: A

Explanation:

The best disk type for a high-performance financial application is a SAS SSD. A SAS SSD (Serial Attached SCSI Solid State Drive) is a type of storage device that uses flash memory chips to store data and has a SAS interface to connect to a server or a storage array. A SAS SSD offers high speed, low latency, high reliability, and high durability compared to other types of disks, such as SATA SSDs, SAS HDDs, or SATA HDDs. A SAS SSD can handle high I/O workloads and deliver consistent performance for applications that require fast data access and processing.

Reference:
https://www.hp.com/us-en/shop/tech-takes/sas-vs-sata

Question: 49

Which of the following DR testing scenarios is described as verbally walking through each step of the DR plan in the context of a meeting?

A. Live failover
B. Simulated failover
C. Asynchronous
D. Tabletop

Answer: D

Explanation:

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

Question: 50

When configuring networking on a VM, which of the following methods would allow multiple VMs to share the same host IP address?

A. Bridged
B. NAT
C. Host only
D. vSwitch

Answer: B

Explanation:

The method that would allow multiple VMs to share the same host IP address is NAT. NAT (Network Address Translation) is a technique that allows multiple devices to use a single public IP address by mapping their private IP addresses to different port numbers. NAT can be used for VM networking to enable multiple VMs on the same host to access the internet or other networks using the host's IP address. NAT can also provide security benefits by hiding the VMs' private IP addresses from external networks.

Reference: https://www.virtualbox.org/manual/ch06.html

Question: 51

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

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

Answer: D

Explanation:

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

device. Downgrading the firmware can help resolve any problems caused by a faulty firmware update and restore normal operation of the device.

Question: 52

A server administrator has noticed that the storage utilization on a file server is growing faster than planned. The

administrator wants to ensure that, in the future, there is a more direct relationship between the number of users using the server and the amount of space that might be used. Which of the following would BEST enable this correlation?

A. Partitioning
B. Deduplication
C. Disk quotas
D. Compression

Answer: C

Explanation:

The best way to ensure that there is a more direct relationship between the number of users using the server and the amount of space that might be used is to implement disk quotas. Disk quotas are a feature that allows a server administrator to limit the amount of disk space that each user or group can use on a file server. Disk quotas can help manage storage utilization, prevent disk space exhaustion, and enforce fair usage policies. Disk quotas can also provide reports and alerts on disk space usage and quota status.

Question: 53

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

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

Answer: D

Explanation:

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

Question: 54

Which of the following can be BEST described as the amount of time a company can afford to be down during recovery from an outage?

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

Answer: C

Explanation:

The term that best describes the amount of time a company can afford to be down during recovery from an outage is RTO. RTO (Recovery Time Objective) is a metric that defines the maximum acceptable downtime for an application, system, or process after a disaster or disruption. RTO helps determine the level of urgency and resources required for restoring normal business operations. RTO is usuallymeasured in minutes, hours, or days, depending on the criticality and impact of the service. Reference:
https://whatis.techtarget.com/definition/recovery-time-objective-RTO

Question: 55

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

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

Answer: C

Explanation:

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

Question: 56

A systems administrator is performing maintenance on 12 Windows servers that are in different racks at a large datacenter. Which of the following would allow the administrator to perform maintenance on all 12 servers without having to physically be at each server? (Choose two.)

A. Remote desktop
B. IP KVM
C. A console connection
D. A virtual administration console
E. Remote drive access
F. A crash cart

Answer: A,B

Explanation:

The methods that would allow the administrator to perform maintenance on all 12 servers without having to physically be at each server are remote desktop and IP KVM. Remote desktop is a feature that allows a user to access and control another computer over a network using a graphical user interface (GUI). Remote desktop can enable remote administration, troubleshooting, and maintenance of servers without requiring physical presence at the server location. IP KVM (Internet Protocol Keyboard Video Mouse) is a device that allows a user to access and control multiple servers over a network using a single keyboard, monitor, and mouse. IP KVM can provide remote access to servers regardless of their operating system or power state, and can also support virtual media and serial console functions.

Reference:
https://www.blackbox.be/en-be/page/27559/Resources/Technical-Resources/Black-Box-Explains/kvm/ Benefits-of-using-KVM-over-IP

Question: 57

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

># Enforcing

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

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

Answer: A

Explanation:

The command that the administrator should issue to configure MySQL successfully is setenforce 0. This command sets the SELinux (Security-Enhanced Linux) mode to permissive, which means that SELinux will not enforce its security policies and will only log any violations. SELinux is a feature that provides mandatory access control (MAC) for Linux systems, which can enhance the security and prevent unauthorized access or modification of files and processes. However, SELinux can also interfere with some applications or services that require specific permissions or ports that are not allowed by SELinux by default. In this case, MySQL may not be able to run properly due to SELinux restrictions. To resolve this issue, the administrator can either disable SELinux temporarily by using setenforce 0, or permanently by editing the /etc/selinux/config file and setting SELINUX=disabled. Alternatively, the administrator can configure SELinux to allow MySQL to run by using commands such as semanage or setsebool.

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

Question: 58

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.

Question: 59

Which of the following server types would benefit MOST from the use of a load balancer?

A. DNS server
B. File server
C. DHCP server

D. Web server

Answer: D

Explanation:

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

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

Question: 60

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

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

Answer: D

Explanation:

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

Question: 61

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

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

Answer: D,E

Explanation:

The most likely reasons why the technician is unable to access a server's package repository internally or externally are that the external firewall is blocking access and that the default gateway is incorrect. A package repository is a source of software packages that can be installed or updated on a server using a package manager tool. A package repository can be accessed over a network using a URL or an IP address. However, if there are any network issues or misconfigurations, the access to

the package repository can be blocked or failed. An external firewall is a device or software that controls the incoming and outgoing network traffic based on predefined rules or policies. An external firewall can block access to a package repository if it does not allow traffic on certain ports or protocols that are used by the package manager tool. A default gateway is a device or address that routes network traffic from one network to another network. A default gateway can be incorrect if it does not match the actual device or address that connects the server's network to other networks, such as the internet. An incorrect default gateway can prevent the server from reaching the package repository over other networks.
QUESTI O N N O: 63 Correct Answer: A The RAID level Explanation:that provides the highest possible capacity for a storage array is RAID 0. RAID 0 is a type of RAID level that provides performance enhancement by using striping. Striping means dividing data into blocks and distributing them across multiple disks to increase speed and capacity. RAID 0 does not provide any fault tolerance or redundancy, as it does not use any parity or mirroring techniques. RAID 0 uses all of the available disk space for data storage, without losing any space for overhead. Therefore, RAID 0 provides the highest possible capacity for a storage array, but also has the highest risk of data loss.

Question: 62

A server administrator was asked to build a storage array with the highest possible capacity. Which of the following RAID levels should the administrator choose?

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

Answer: A

Explanation:

The RAID level that provides the highest possible capacity for a storage array is RAID 0. RAID 0 is a type of RAID level that provides performance enhancement by using striping. Striping means dividing data into blocks and distributing them across multiple disks to increase speed and capacity. RAID 0 does not provide any fault tolerance or redundancy, as it does not use any parity or mirroring techniques. RAID 0 uses all of the available disk space for data storage, without losing any space for overhead. Therefore, RAID 0 provides the highest possible capacity for a storage array, but also has the highest risk of data loss.

Reference: https://www.thinkmate.com/inside/articles/what-is-raid

Question: 63

A technician needs to deploy an operating system that would optimize server resources. Which of the following server installation methods would BEST meet this requirement?

A. Full
B. Bare metal
C. Core
D. GUI

Answer: C

Explanation:

The server installation method that would optimize server resources is core. Core is a minimal installation option that is available for some operating systems, such as Windows Server and Linux. Core installs only the essential components and features of the operating system, without any graphical user interface (GUI) or other unnecessary services or applications. Core reduces the disk footprint, memory usage, CPU consumption, and attack surface of the server, making it more efficient and secure. Core can be managed remotely using command-line tools, PowerShell, or GUI tools.

Reference:
https://docs.microsoft.com/en-us/windows-server/administration/performance-tuning/hardware/

Question: 64

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

A. Disable port 389 on the server

B. Move traffic from port 389 to port 443
C. Move traffic from port 389 to port 637
D. Enable port 389 for web traffic

Answer: A

Explanation:

The best way to complete the request to harden the server is to disable port 389 on the server. Port 389 is the default port used by LDAP (Lightweight Directory Access Protocol), which is a protocol that allows access and modification of directory services over a network. LDAP can be used for authentication, authorization, or information retrieval purposes. However, LDAP does not encrypt its data by default, which can expose sensitive information or credentials to attackers who can intercept or modify the network traffic. Therefore, port 389 should be disabled on a web server that only hosts websites and does not need LDAP functionality. Alternatively, port 636 can be used instead of port 389 to enable LDAPS (LDAP over SSL/TLS), which encrypts the data using SSL/TLS certificates.

Question: 65

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.

Question: 66

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

A. Restart the server

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

Answer: C

Explanation:

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

Question: 67

Which of the following would MOST likely be part of the user authentication process when implementing SAML across multiple applications?

A. SSO
B. LDAP
C. TACACS
D. MFA

Answer: A

Explanation:

The term that is most likely part of the user authentication process when implementing SAML across multiple applications is SSO. SSO (Single Sign-On) is a way for users to be authenticated for multiple applications and services at once. With SSO, a user signs in at a single login screen and can then use a number of apps without having to enter their credentials again. SSO improves user experience and security by reducing password fatigue and phishing risks. SAML (Security Assertion Markup Language) is a protocol that enables SSO by providing a standardized way to exchange authentication and authorization data between an identity provider (IdP) and a service provider (SP). SAML uses XML-based messages called assertions to communicate user identity and attributes between parties. Reference:

https://www.onelogin.com/learn/how-single-sign-on-works

Question: 68

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

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

Answer: B

Explanation:

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

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

Question: 69

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

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

Answer: C

Explanation:

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

Question: 70

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

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

Answer: B

Explanation:

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

Question: 71

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

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

Answer: A

Explanation:

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

Question: 72

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

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

Answer: B

Explanation:

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

Question: 73

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.

Question: 74

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

\\server1\ITDept\

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

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

Answer: C

Explanation:

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

Question: 75

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

A. the correct firewall zone is active

B. the latest firmware was applied
C. NTP is running on the database system
D. the correct dependencies are installed

Answer: C

Explanation:

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

Question: 76

A technician is connecting a Linux server to a share on a NAS. Which of the following is the MOST

appropriate native protocol to use for this task?

A. CIFS B. FTP
C. SFTP D. NFS

Answer: D

Explanation:

The most appropriate native protocol to use for connecting a Linux server to a share on a NAS is NFS. NFS (Network File System) is a protocol that allows file sharing and remote access over a network. NFS is designed for Unix-like operating systems, suchas Linux, and supports features such as symbolic links, hard links, file locking, and file permissions. NFS uses mount points to attach remote file systems to local file systems, making them appear as if they are part of the local file system. NFS can provide fast and reliable access to files stored on a NAS (Network Attached Storage), which is a device that provides centralized storage for network devices.

Question: 77

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

Question: 78

A server is reporting a hard drive S.M.A.R.T. error. When a technician checks on the drive, however, it appears that all drives in the server are functioning normally. Which of the following is the reason for this issue?

A. A S.M.A.R.T. error is a predictive failure notice. The drive will fail in the near future and should be replaced at the next earliest time possible
B. A S.M.A.R.T. error is a write operation error. It has detected that the write sent to the drive was incorrectly formatted and has requested a retransmission of the write from the controller
C. A S.M.A.R.T. error is simply a bad sector. The drive has marked the sector as bad and will continue to function properly
D. A S.M.A.R.T. error is an ECC error. Due to error checking and correcting, the drive has corrected the missing bit and completed the write operation correctly.

Answer: A

Explanation:

A S.M.A.R.T. error is a predictive failure notice. The drive will fail in the near future and should be replaced at the next earliest time possible. S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology) is a feature that monitors the health and performance of hard drives and alerts the user of any potential problems or failures. S.M.A.R.T. can detect various indicators of drive degradation, such as bad sectors, read/write errors, temperature, or spin-up time. If a S.M.A.R.T. error is reported, it means that the drive has exceeded a predefined threshold of acceptable operation and is likely to fail soon. The drive may still function normally for a while, but it is recommended to back up the data and replace the drive as soon as possible to avoid data loss or system downtime.

Question: 79

A server administrator has been creating new VMs one by one. The administrator notices the system requirements are very similar, even with different applications. Which of the following would help the administrator accomplish this task in the SHORTEST amount of time and meet the system requirements?

A. Snapshot

B. Deduplication

C. System Restore

D. Template

Answer: D

Explanation:

The method that would help the administrator accomplish the task of creating new VMs in the shortest amount of time and meet the system requirements is template. A template is a preconfigured virtual machine image that contains an operating system, applications, settings, and other components. A template can be used to create multiple identical or customized VMs quickly and easily, without having to install and configure each VM from scratch. A template can save time and ensure consistency across VMs.

Question: 80

Which of the following steps in the troubleshooting theory should be performed after a solution has been implemented? (Choose two.)

A. Perform a root cause analysis

B. Develop a plan of action

C. Document the findings

D. Escalate the issue

E. Scope the issue

F. Notify the users

Answer: C,F

Explanation:

The steps in the troubleshooting theory that should be performed after a solution has been implemented are document the findings and notify the users. The troubleshooting theory is a systematic process of identifying and resolving problems or issues with a system or device. The troubleshooting theory consists of several steps that can be summarized as follows: Identify the problem: Gather information, scope the issue, establish a theory of probable cause. Establish a plan of action: Test the theory, determine next steps, escalate if necessary. Implement the solution: Execute the plan, verify functionality, prevent recurrence.
Document the findings: Record actions taken, outcomes achieved, lessons learned. Notify the users: Communicate resolution status, confirm satisfaction, provide follow-up. Documenting the findings is an important step that helps create a record of what was done and why, what worked and what didn't, and what can be improved or avoided in the future. Documenting the

findings can also help with reporting, auditing, compliance, or training purposes. Notifying the users is another important step that helps inform the affected parties of what was done and how it was resolved, confirm that the problem is fixed and that they are satisfied with the outcome, and provide any follow-up instructions or recommendations.

Question: 81

Which of the following allows for a connection of devices to both sides inside of a blade enclosure?

A. Midplane
B. Active backplane
C. Passive backplane
D. Management module

Answer: A

Explanation:

The component that allows for a connection of devices to both sides inside of a blade enclosure is midplane. A midplane is a board or panel that connects two sets of connectors or devices in parallel with each other. A midplane is typically used in blade enclosures or chassis to provide power and data connections between blade servers on one side and power supplies, cooling fans, switches, or management modules on the other side. A midplane can also act as a backplane by providing bus signals or communication channels between devices.

Question: 82

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

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

Answer: B

Explanation:

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

Question: 83

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

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

Answer: D

Explanation:

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

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

The most efficient way to perform the task of deploying five VMs with the same type of configuration is to

use a VM template. A template is a preconfigured virtual machine image that contains an operating system, applications, settings, and other components. A template can be used to create multiple identical or customized VMs quickly and easily, without having to install and configure each VM from scratch. A template can save time and ensure consistency across VMs.

Question: 84

A global organization keeps personnel application servers that are local to each country. However, a security audit shows these application servers are accessible from sites in other countries. Which of the following hardening techniques should the organization use to restrict access to only sites that are in the same country?

A. Configure a firewall
B. Close the unneeded ports
C. Install a HIDS
D. Disable unneeded services.

Answer: A

Explanation:

Monitors Network Traffic

Reference:https://www.fortinet.com/resources/cyberglossary/benefits-of-firewall

Question: 85

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

A. RFID
B. Proximity readers
C. Signal blocking
D. Camouflage
E. Reflective glass
F. Bollards

Answer: C,E

Explanation:

The best solutions to resolve the concern of transmissions from the building being detected from outside are signal blocking and reflective glass. Signal blocking is a method of preventing or

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

Question: 86

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

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

Question: 87

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

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

Answer: D

Explanation:

Few people are expected to use the database at the same time and users don't need to customize the design of the database.

Reference:https://support.microsoft.com/en-us/office/ways-to-share-an-access-desktop-database- 03822632-da43-4d8f-ba2a-68da245a0446

The server role that the administrator will most likely need to install for a server that will hold large amounts of information and will need to be accessed by multiple users at the same time is database. A database is a collection of structured data that can be stored, queried, manipulated, and analyzed using various methods and tools. A database server is a server that hosts one or more databases and provides access to them over a network. A database server can handle large amounts of information and support concurrent requests from multiple users or applications.

Question: 88

Users at a company work with highly sensitive dat
a. The security department implemented an administrative and technical control to enforce least-

privilege access assigned to files. However, the security department has discovered unauthorized data exfiltration. Which of the following is the BEST way to protect the data from leaking?

A. Utilize privacy screens.
B. Implement disk quotas.
C. Install a DLP solution.
D. Enforce the lock-screen feature.

Answer: C

Explanation:

Components of a Data Loss Solution

Reference:https://www.imperva.com/learn/data-security/data-loss-prevention-dlp/

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

Question: 89

A server administrator needs to create a new folder on a file server that only specific users can access. Which of the following BEST describes how the server administrator can accomplish this task?

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

Answer: C

Explanation:

The top portion of the dialog box lists the users and/or groups that have access to the file or folder.

Reference:https://www.uwec.edu/kb/article/drives-establishing-windows-file-and-folder-level- permissions/

Question: 90

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

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

Answer: B

Explanation:

Unexpected shutdowns. If the system is randomly shutting down or rebooting, the most likely cause is a heat problem.

Reference:https://www.microsoftpressstore.com/articles/article.aspx?p=2224043&seqNum=3

Question: 91

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

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

Answer: B

Explanation:

Secure Sockets Layer (SSL): SSL and its successor Transport Layer Security (TLS) enable client and server computers to establish a secure connection session and manage encryption and decryption activities.

Reference:https://paginas.fe.up.pt/~als/mis10e/ch8/chpt8-4bullettext.htm

Question: 92

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

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

Answer: A

Explanation:

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

Question: 93

HOTSPOT

A systems administrator deployed a new web proxy server onto the network. The proxy server has two interfaces: the first is connected to an internal corporate firewall, and the second is connected to an internet-facing firewall. Many users at the company are reporting they are unable to access the Internet since the new proxy was introduced. Analyze the network diagram and the proxy server's host routing table to resolve the Internet connectivity issues.

INSTRUCTIONS

Perform the following steps:

1. Click on the proxy server to display its routing table.
2. Modify the appropriate route entries to resolve the Internet connectivity issue.

If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

Proxy Server Routing Table

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

Answer:

Explanation:

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

Question: 94

A systems administrator needs to configure a new server and external storage for a new production application environment. Based on end-user specifications, the new solution needs to adhere to the following basic requirements:

1. The OS must be installed in a separate disk partition. In case of hard drive failure, it cannot be affected.
2. Application data IOPS performance is a must.

3. Data availability is a high priority, even in the case of multiple hard drive failures.

Which of the following are the BEST options to comply with the user requirements? (Choose three.)

A. Install the OS on a RAID 0 array.
B. Install the OS on a RAID 1 array.
C. Configure RAID 1 for the application data.
D. Configure RAID 5 for the application data.
E. Use SSD hard drives for the data application array.
F. Use SATA hard drives for the data application array.
G. Use a single JBOD for OS and application data.

Answer: B,D,E

Explanation:

To comply with the user requirements, the best options are to install the OS on a RAID 1 array, configure RAID 5 for the application data, and use SSD hard drives for the data application array. Here is why:
RAID 1 is a mirroring technique that creates an exact copy of data on two disks. This provides redundancy and fault tolerance in case of hard drive failure. RAID 1 also improves read performance since either disk can be read at the same time. Therefore, installing the OS on a RAID 1 array meets the first requirement of separating the OS from the application data and protecting it from hard drive failure.
RAID 5 is a striping technique with parity that distributes data and parity blocks across three or more disks. This provides improved performance and storage efficiency compared to RAID 1, as well as fault tolerance in case of a single disk failure. Therefore, configuring RAID 5 for the application data meets the second and third requirements of providing high IOPS performance and data availability. SSD hard drives are solid-state drives that use flash memory to store data. They have no moving parts and offer faster read and write speeds, lower latency, and lower power consumption than traditional HDDs. Therefore, using SSD hard drives for the data application array meets the second requirement of providing high IOPS performance.
Reference:
https://phoenixnap.com/kb/raid-levels-and-types https://en.wikipedia.org/wiki/Standard_RAID_levels

Question: 95

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

```
ipconfig /all

IPv4 address: 192.168.1.5
```

```
Subnet mask: 255.255.255.0
Default gateway: 192.168.1.1

pinging 192.168.1.1 with 32 bytes of data:

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

Which of the following caused the issue?

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

Answer: A

Explanation:

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

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

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

destinations, not a reply from a different IP address.

Reference:

https://askleo.com/what_is_ping_and_what_does_its_output_tell_me/
https://learn.microsoft.com/en-us/windows-server/administration/windows-commands/ping

Question: 96

A server administrator is swapping out the GPU card inside a server. Which of the following actions should the administrator take FIRST?

A. Inspect the GPU that is being installed.
B. Ensure the GPU meets HCL guidelines. C. Shut down the server.
D. Disconnect the power from the rack.

Answer: C

Explanation:

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

Question: 97

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

Question: 98

Which of the following relates to how much data loss a company agrees to tolerate in the event of a disaster?

A. RTO
B. MTBF
C. PRO
D. MTTR

Answer: A

Explanation:

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

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

Question: 99

A server administrator is testing a disaster recovery plan. The test involves creating a downtime scenario and taking the necessary steps. Which of the following testing methods is the administrator MOST likely performing?

A. Backup recovery
B. Simulated
C. Tabletop
D. Live failover

Answer: D

Explanation:

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

Question: 100

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

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

Answer: C

Explanation:

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

Question: 101

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 rebuild ing. 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/

Question: 102

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

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

Answer: C

Explanation:

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

Question: 103

A remote physical server is unable to communicate to the network through the available NICs, which were misconfigured. However, the server administrator is still able to configure theserver remotely.
Which of the following connection types is the server administrator using to access the server?

A. Out-of-band management
B. Crash cart access
C. Virtual administrator console
D. Local KVM setup
E. RDP connection

Answer: A

Explanation:

The connection type that the server administrator is using to access the server remotely is out-ofband management. Out-of-band management is a method of accessing and controlling a server through a dedicated network interface or port that is separate from the regular data network. Out- of-band management allows administrators to perform tasks such as rebooting, configuring, troubleshooting, or updating a server even if the server is ofline or unresponsive through the regular network. Out-of-band management can use protocols such as IPMI, iLO, DRAC, or BMC.Reference:https://www.ibm.com/cloud/learn/out-of-band-management

Question: 104

A system administrator has been alerted to a zero-day vulnerability that is impacting a service enabled on a server OS. Which of the following would work BEST to limit an attacker from exploiting this vulnerability?

A. Installing the latest patches
B. Closing open ports
C. Enabling antivirus protection
D. Enabling a NIDS

Answer: A

Explanation:

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

Question: 105

A server administrator has connected a new server to the network. During testing, the administrator discovers the server is not reachable via server but can be accessed by IP address. Which of the following steps should the server administrator take NEXT? (Select TWO).

A. Check the default gateway.
B. Check the route tables.
C. Check the hosts file.
D. Check the DNS server.
E. Run the ping command.
F. Run the tracert command

Answer: C,D

Explanation:

If the server is not reachable by name but can be accessed by IP address, it means that there is a problem with name resolution. The hosts file and the DNS server are both responsible for mapping hostnames to IP addresses. Therefore, the server administrator should check these two files for any errors or inconsistencies that might prevent the server from being resolved by name. Reference: https://www.howtogeek.com/662249/how-to-edit-the-hosts-file-on-linux/ https://www.howtogeek.com/164981/how-to-use-nslookup-to-check-domain-name-information-in-microsoft-windows/

Question: 106

An administrator needs to disable root login over SSH. Which of the following tiles should be edited to complete this task?

A. /root.ssh/sshd/config
B. /etc.ssh/sshd_config
C. /root/.ssh/ssh_config
D. /etc.sshs_shd_config

Answer: B

Explanation:

To disable root login over SSH, the server administrator needs to edit the SSH configuration file located at /etc/ssh/sshd_config. This file contains various settings for the SSH daemon that runs on the server and accepts incoming SSH connections. The administrator needs to find the line that says

PermitRootLogin and change it to no or comment it out with a # symbol. Then, the administrator needs to restart the SSH service for the changes to take effect.Reference:https://www.howtogeek.com/828538/how-and-why-to-disable-root-login-over- ssh-on-linux/

Question: 107

Users have noticed a server is performing below Baseline expectations. While diagnosing me server, an administrator discovers disk drive performance has degraded. The administrator checks the diagnostics on the RAID controller and sees the battery on me controller has gone bad. Which of the following is causing the poor performance on the RAID array?

A. The controller has disabled the write cache.
B. The controller cannot use all the available channels.
C. The drive array is corrupt.
D. The controller has lost its configuration.

Answer: A

Explanation:

The write cache is a feature of some RAID controllers that allows them to temporarily store data in a fast memory buffer before writing it to the disk drives. This improves the performance and efficiency of write operations, especially for random and small writes. However, if the battery on the controller goes bad, the controller may disable the write cache to prevent data loss in case of a power failure. This can degrade the disk drive performance significantly, as every write operation will have to wait for the disk drives to complete. Reference: https://www.dell.com/support/kbdoc/en-us/000131486/understanding-raid-controller-battery-learn- cyclehttps://www.techrepublic.com/article/understanding-

raid-controller-write-cache/

Question: 108

A server technician notices a server is very low on disk space. Upon inspecting the disk utilization, the technician discovers server logs are taxing up a large amount of space. There is no central log server. Which of the following would help free up disk space?

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

Answer: B

Explanation:

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

Question: 109

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

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

Answer: A,C

Explanation:

The power supply and the hard drive are two components that can most likely be swapped out while the server is still running, if they support hot swapping or hot plugging. Hot swapping or hot plugging means that the device can be added or

removed without shutting down the system. The operating system automatically recognizes the changes that have been made. This feature is useful for minimizing downtime and improving availability. The CPU, the GPU, the cache, and the RAM are not hot swappable and require the system to be powered off before replacing them. Reference: https://www.geeksforgeeks.org/what-is-hot-swapping/https://www.howtogeek.com/268249/what- is-hot-swapping-and-what-devices-support-it/

Question: 110

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

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

Answer: B

Explanation:

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

Question: 111

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

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

Answer: B

Explanation:

If users cannot access a new server by name, but the server does respond to a ping request using its IP address, it means that there is a problem withname resolution. The DNS (Domain Name System) is a service that maps hostnames to IP addresses and vice versa. Therefore, the best step to perform next is to examine the DNS to see if the new server record exists and matches its IP address. If not, the DNS record needs to be added or updated accordingly. Running the tracert command from a workstation would not help with name resolution, as it only shows the route taken by packets to reach a destination by IP address. Correcting the missing DHCP scope would not help either, as DHCP (Dynamic Host Configuration Protocol) only assigns IP addresses and other network settings to clients, but does not resolve names. Updating the workstation hosts file would be a temporary workaround, but not a permanent solution, as it would require manually editing every workstation's hosts file with the new server's name and IP address. Reference:

https://www.howtogeek.com/164981/how-to-use-nslookup-to-check-domain-name-information-in- microsoft-windows/https://www.howtogeek.com/howto/27350/beginner-geek-how-to-edit-your- hosts-file/

Question: 112

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

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

Answer: D

Explanation:

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

Question: 113

An administrator is investigating a physical server mat will not Boot into the OS. The server has three hard drives configured

in a RAID 5 array. The server passes POST, out the OS does not load. The administrator verities the CPU and RAM are Doth seated correctly and checks the dual power supplies. The administrator then verifies all the BIOS settings are correct and connects a bootable USB drive in the server, and the OS loads correctly. Which of the following is causing the issue?

A. The page file is too small.
B. The CPU has failed.
C. There are multiple failed hard drives.
D. There are mismatched RAM modules.
E. RAID 5 requires four drives

Answer: C

Explanation:

If a server has three hard drives configured in a RAID 5 array, it means that the data is striped across all three drives with parity information. RAID 5 can tolerate one drive failure without losing data, but not two or more. If there are multiple failed hard drives, the RAID 5 array will become corrupted and the OS will not load. The other options are not likely to cause the issue, as the server passes POST, the CPU and RAM are seated correctly, the BIOS settings are correct, and the OS loads from a bootable USB drive. RAID 5 does not require four drives, it can work with three or more.Reference:https://www.technewstoday.com/what-is-a-raid-5/

Question: 114

Which of the following BEST measures now much downtime an organization can tolerate Curing an unplanned outage?

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

Answer: C

Explanation:

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

Question: 115

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. Reference: 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/

Question: 116

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

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

Answer: B

Explanation:

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

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

Question: 117

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

A. LVM
B. DiskPart
C. fdisk
D. Format

Answer: A

Explanation:

LVM (Logical Volume Manager) is a tool that allows the technician to increase a partition's size in the future without reformatting it on a Linux server. LVM creates logical volumes that can span across multiple physical disks or partitions and can be resized dynamically without losing data. LVM also provides other features such as snapshots, encryption, and RAID. DiskPart, fdisk, and Format are tools that can be used to partition and format disks, but they do not allow increasing a partition's size without reformatting it. Reference: https://www.howtogeek.com/howto/40702/how-to- manage-and-use-lvm-logical-volume-management-in-ubuntu/ https://www.howtogeek.com/school/using-windows-admin-tools-like-a-pro/lesson2/https://www.howtogeek.com/howto/17001/how-to-format-a-usb-drive-in-ubuntu- using-gparted/

Question: 118

A very old PC is running a critical, proprietary application in MS-DOS. Administrators are concerned about the stability of this computer. Installation media has been lost, and the vendor is out of business. Which of the following would be the BEST course of action to preserve business continuity?

A. Perform scheduled chkdsk tests.
B. Purchase matching hardware and clone the disk.
C. Upgrade the hard disk to SSD.
D. Perform quarterly backups.

Answer: B

Explanation:

The best course of action to preserve business continuity for a very old PC that is running a critical, proprietary application in MS-DOS is to purchase matching hardware and clone the disk. This way, the technician can create an exact copy of the PC's

configuration and data on another PC that has the same specifications and compatibility. This will ensure that the application can run smoothly on the new PC without any installation or configuration issues. Performing scheduled chkdsk tests would not help, as chkdsk is a tool that checks and repairs disk errors, but does not prevent hardware failures or software compatibility issues. Upgrading the hard disk to SSD would not help either, as SSDs may not be compatible with the old PC or the MS-DOS operating system. Performing quarterly backups would help with data protection, but not with hardware availability or software compatibility. Reference: https://www.howtogeek.com/199068/how-to-upgrade-your-existing-hard-drive-in-under-an-hour/https://www.howtogeek.com/66776/how-to-repair-disk-errors-in-windows- 7/

Question: 119

A technician is attempting to reboot a remote physical Linux server. However, attempts to command a shutdown -- now result in the loss of the SSH connection. The server still responds to pings. Which of the following should the technician use to command a remote shutdown?

A. virtual serial console
B. A KVM
C. An IDRAC
D. A crash cart

Answer: C

Explanation:

An IDRAC (Integrated Dell Remote Access Controller) is a tool that can be used to command a remote shutdown of a physical Linux server. An IDRAC is a hardware device that provides out-of-band management for Dell servers. It allows the technician to access the server's console, power cycle, reboot, or shut down the server remotely using a web interface or a command-line interface. An IDRAC does not depend on the operating system or network connectivity of the server. A virtual serial console is a tool that can be used to access a remote virtual machine's console using a serial port connection. A KVM (Keyboard, Video, Mouse) switch is a device that allows the technician to switch between different computer sources using the same keyboard, monitor, and mouse. A crash cart is a mobile unit that contains a keyboard, monitor, mouse, and other tools that can be connected to a physical server for troubleshooting purposes. Reference: https://www.dell.com/support/kbdoc/en-us/000131486/understanding-the-idrac
https://www.howtogeek.com/799968/what-is-a-kvm-

switch/https://www.techopedia.com/definition/1032/business-impact-analysis-bia

Question: 120

An administrator is troubleshooting a RAID issue in a failed server. The server reported a drive failure, and then it crashed and would no longer boot. There are two arrays on the failed server: a two-drive RAIO 0 set tor the OS, and an eight-drive RAID 10 set for dat a. Which of the following failure scenarios MOST likely occurred?

A. A drive failed in the OS array.
B. A drive failed and then recovered in the data array.
C. A drive failed in both of the arrays.
D. A drive failed in the data array.

Answer: A

Explanation:

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

Question: 121

Which of the following technologies would allow an administrator to build a software RAID on a Windows server?

A. Logical volume management
B. Dynamic disk
C. GPT
D. UEFI

Answer: B

Explanation:

Dynamic disk is a technology that allows an administrator to build a software RAID on a Windows server. Dynamic disk is a type of disk management that supports creating volumes that span multiple

disks, stripe data across disks, mirror data between disks, or use parity for fault tolerance. Dynamic disk can be used to create RAID 0 (striping), RAID 1 (mirroring), RAID 5 (striping with parity), or spanned volumes on Windows servers. Logical volume management is a technology that allows creating and resizing logical volumes on Linux servers. GPT (GUID Partition Table) is a standard for defining the partition structure on a disk. UEFI (Unified Extensible Firmware Interface) is a specification for the interface between the operating system and the firmware. Reference:
https://www.howtogeek.com/school/using-windows-admin-tools-like-a-pro/lesson2/
https://www.howtogeek.com/howto/40702/how-to-manage-and-use-lvm-logical-volume- management-in-ubuntu/
https://www.howtogeek.com/193669/whats-the-difference-between-gpt- and-mbr-when-partitioning-a-drive/https://www.howtogeek.com/56958/htg-explains-how-uefi-will- replace-the-bios/

Question: 122

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

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

Answer: A

Explanation:

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

Question: 123

A datacenter in a remote location lost power. The power has since been restored, but one ol the servers has not come back online. After some investigation, the server is found to still be powered off. Which of the following is the BEST method to power on the server remotely?

A. Crash cart
B. Out-of-band console
C. IP KVM
D. RDP

Answer: B

Explanation:

Out-of-band console is a tool that can be used to command a remote shutdown of a physical Linux server. Out-of-band console is a method of accessing a server's console through a dedicated management port or device that does not rely on the server's operating system or network connection. Out-of-band console can be used to power cycle, reboot, update firmware, monitor performance, and perform other tasks remotely even if the server is unresponsive or ofline. Crash cart is a mobile unit that contains a keyboard, monitor, mouse, and other tools that can be used to troubleshoot a server on-site, but it requires physical access to the server. IP KVM (Internet Protocol Keyboard Video Mouse) switch is a hardware device that allows remote access to multiple servers using a web browser or a client software,

but it requires network connectivity and may not work if the SSH connection is lost. RDP (Remote Desktop Protocol) is a protocol that allows remote access to a Windows server's graphical user interface, but it does not work on Linux servers and requires network connectivity. Reference: https://www.techopedia.com/definition/13623/crash-cart https://www.techopedia.com/definition/13624/kvm- switchhttps://www.techopedia.com/definition/3422/remote-desktop-protocol-rdp

Question: 124

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

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

Answer: C

Explanation:

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

Question: 125

A backup application is copying only changed files each line 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 synthetize 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. Reference: https://www.nakivo.com/blog/what-is-synthetic- backup/https://www.howtogeek.com/192115/what-you-need-to-know-about-creating-system- image-backups/

Question: 126

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

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

Answer: D

Explanation:

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

Question: 127

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

A. Wire the four network cables and the two power cables through the cable management arm using appropriate-length cables.

B. Run the tour network cables up the left side of the rack to the top of the rack switch. Run the two power cables down the right side of the rack toward the UPS.

C. Use the longest cables possible to allow for adjustment of the server rail within the rack.

D. Install an Ethernet patch panel and a PDU to accommodate the network and power cables.

Answer: B

Explanation:

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

Question: 128

An administrator is configuring a host-based firewall tor 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: A,B,H

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. Reference: 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/

Question: 129

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

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

Answer: E,F

Explanation:

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

Question: 130

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

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

Answer: A

Explanation:

The first thing that the technician should do before installing a new server OS on legacy server hardware is to consult the HCL (Hardware Compatibility List) to ensure everything is supported. The HCL is a list of hardware devices and components that are tested and certified to work with a specific

OS or software product. The HCL helps to avoid compatibility issues and performance problems that may arise from using unsupported or incompatible hardware. Migrating the physical server to a virtual server may be a good option to improve scalability and flexibility, but it requires additional hardware and software resources and may not be feasible for legacy server hardware. Low-level formatting the hard drives may be a good practice to erase any old data and prepare the drives for a new OS installation, but it does not guarantee that the hardware will work with the new OS. Making sure the case and the fans are free from dust may be a good practice to ensure proper cooling and prevent overheating, but it does not guarantee that the hardware will work with the new OS. Reference: https://www.howtogeek.com/190014/virtualization-basics-understanding-techniques- and-fundamentals/ https://www.howtogeek.com/173353/how-to-low-level-format-or-write-zeros- to-a-hard-drive/https://www.howtogeek.com/303282/how-to-manage-your-pcs-fans-for-optimal- airflow-and-cooling/

Question: 131

Which of me following BEST describes a disaster recovery site with a target storage array that receives replication traffic and servers that are only powered on In the event of a disaster?

A. Cloud
B. Cold
C. Hot
D. Warm

Answer: D

Explanation:

A warm site is a type of disaster recovery site that has a target storage array that receives replication traffic and servers that are only powered on in the event of a disaster. A warm site is a compromise between a hot site and a cold site. A warm site has some equipment and data ready, but requires some configuration and restoration before resuming operations. A warm site is usually located in a different geographic area than the primary site and has redundant power, cooling, network, and security systems. A warm site is suitable for organizations that can tolerate some downtime and data loss in case of a disaster. A cloud site is a type of disaster recovery site that uses cloud-based resources and platforms to store backups and restore data and applications after a disaster. A cold site is a type of disaster recovery site that has only basic infrastructure and space available, but requires significant setup and installation before resuming operations. A hot site is a type of disaster recovery site that has all the equipment and data ready to resume operations as soon as possible after a disaster. Reference: https://www.techopedia.com/definition/11172/hot-site https://www.techopedia.com/definition/11173/warm-site https://www.techopedia.com/definition/11174/cold-sitehttps://www.techopedia.com/definition/29836/cloud-recovery

Question: 132

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

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

Answer: B

Explanation:

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

Question: 133

Which of me following should be placed at the top of a Bash script to ensure it can be executed?

A. bash
B. !execute
C. #!
D. @eacho off

Answer: C

Explanation:

#! is the symbol that should be placed at the top of a Bash script to ensure it can be executed. #! is also known as shebang or hashbang. It is a special notation that tells the operating system which

interpreter to use to run the script. The shebang is followed by the path to the interpreter, such as /bin/bash for Bash,

/bin/python for Python, or /bin/perl for Perl. For example, a Bash script that prints “Hello World” would start with: #!/bin/bash echo “Hello World”

The shebang must be the first line of the script and must not have any spaces between the # and ! symbols. bash is not a valid shebang by itself, as it does not specify the path to the interpreter. !execute is not a valid shebang at all, as it does not start with #. @echo off is a command that disables the echoing of commands in a batch file on Windows, but it has nothing to do with Bash scripts on Linux. Reference: https://www.howtogeek.com/67469/the-beginners-guide-to-shell- scripting-the-basics/ https://www.howtogeek.com/435903/what-is-a-shebang-line/

Question: 134

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

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

Answer: A

Explanation:

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

Question: 135

A Linux administrator created a script that will run at startup. After successfully writing the script, the administrator received me following output when trying to execute the script:

Bash ./startup,sh:Permission denied

Which of the following commands would BEST resolve the error message?

A. Chmod +w startup.sh
B. Chmod 444 startup.sh
C. Chmod+x startup.sh
D. Chmod 466 startUp,sh

Answer: C

Explanation:

This is the command that would best resolve the error message "Bash ./startup.sh: Permission denied" when trying to execute a script on Linux. Chmod is a command that changes the permissions of files or directories on Linux. +x is an option that adds the execute permission to the file or directory for the owner, group, and others. startup.sh is the name of the script file that needs to be executed. By running chmod +x startup.sh, the technician grants execute permission to the script file and allows it to be run by any user. Chmod +w startup.sh would add write permission to the file, but not execute permission. Chmod 444 startup.sh would set read-only permission for all users, but not execute permission. Chmod 466 startup.sh would set read and write permission for the owner and write-only permission forgroup and others, but not execute permission.Reference:https://www.howtogeek.com/437958/how-to-use-the-chmod-command-on-linux

Question: 136

A technician is checking a server rack. Upon entering the room, the technician notices the tans on a particular server in the rack are running at high speeds. This is the only server in the rack that is experiencing this behavior. The ambient temperature in the room appears to be normal. Which of the following is the MOST likely reason why the fans in that server are operating at full speed?

A. The server is In the process of shutting down, so fan speed operations have been defaulted to high.
B. An incorrect fan size was inserted into the server, and the server has had to Increase the fan speed to compensate.
C. A fan failure has occurred, and the other fans have increased speed to compensate.
D. The server is utilizing more memory than the other servers, so it has increased the fans to compensate.

Answer: C

Explanation:

This is the most likely reason why the fans in that server are operating at full speed while the ambient temperature in the room is normal and the other servers in the rack are not experiencing this behavior. A fan failure is a situation where one or more fans in a server stop working or malfunction due to wear and tear, dust, or other factors. This can cause overheating and performance issues on the server. To prevent this, most servers have a fan redundancy feature that allows the other fans to increase their speed and airflow to compensate for the failed fan and maintain a safe temperature level. The server is not likely to be in the process of shutting down, as this would not cause the fans to run at high speeds. An incorrect fan size is not likely to be inserted into the server, as most fans are standardized and compatible with the server chassis and motherboard. The server is not likely to be utilizing more memory than the other servers, as this would not cause a significant increase in temperature or fan speed. Reference: https://www.howtogeek.com/303282/how-to-manage-your-pcs-fans-for-optimal-airflow-and-cooling/https://www.howtogeek.com/174288/how-to-tell-if-your-computer-is-overheating-and- what-to-do-about-it/

Question: 137

Which of me following is the BEST action to perform before applying patches to one of the hosts in a high availability cluster?

A. Disable the heartbeat network.
B. Fallback cluster services.
C. Set the cluster to active-active.
D. Failover all VMs.

Answer: D

Explanation:

This is the best action to perform before applying patches to one of the hosts in a high availability cluster. A high availability cluster is a group of hosts thatact like a single system and provide continuous uptime. A high availability cluster is often used for load balancing, backup, and failover purposes. Failover is a process of transferring workloads from one host to another in case of a failure or maintenance. By failing over all VMs (Virtual Machines) from the host that needs to be patched to another host in the cluster, the technician can ensure that there is no downtime or data loss during the patching process. Disabling the heartbeat network is not a good action to perform, as this would disrupt the communication and synchronization between the hosts in the cluster. Fallback cluster services is not a valid term, but it may refer to restoring cluster services after a failover, which is not

relevant before applying patches. Setting the cluster to active-active is not a good action to perform, as this would increase the load on both hosts and reduce redundancy. Reference: https://www.howtogeek.com/190014/virtualization-basics-understanding-techniques-and-fundamentals/https://www.howtogeek.com/428483/what-is-end-to-end-encryption-and-why-does- it-matter/

Question: 138

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

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

Answer: B

Explanation:

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

Question: 139

A technician has several possible solutions to a reported server issue. Which of the following BEST represents how the technician should proceed with troubleshooting?

A. Determine whether there is a common element in the symptoms causing multiple problems.
B. Perform a root cause analysis.
C. Make one change at a time and test.
D. Document the findings, actions, and outcomes throughout the process.

Answer: C

Explanation:

This is the best way to proceed with troubleshooting when the technician has several possible solutions to a reported server issue. Making one change at a time and testing allows the technician to isolate the cause and effect of each solution and determine which one works best. It also helps to avoid introducing new problems or complicating existing ones by making multiple changes at once. Determining whether there is a common element in the symptoms causing multiple problems is a good step to perform before identifying possible solutions, but not after. Performing a root cause analysis is a good step to perform after resolving the issue, but not during. Documenting the findings, actions, and outcomes throughout the process is a good practice to follow at every step of troubleshooting, but not a specific way to proceed with testing possible

solutions. Reference: https://www.howtogeek.com/174288/how-to-tell-if-your-computer-is-overheating-and-what-to-do-about-it/https://www.howtogeek.com/428483/what-is-end-to-end-encryption-and-why-does-it- matter/

Question: 140

Winch 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 bandwidthconsumption 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. Reference: 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/

Question: 141

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

A. Virtual
B. Blade
C. Tower

D. Rack mount

Answer: C

Explanation:

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

Question: 142

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

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

Answer: D

Explanation:

QSFP+ is a type of connector that should be used to connect a server to a switch that uses 40G network connectivity. QSFP+ (Quad Small Form-factor Pluggable Plus) is a compact, hot-pluggable transceiver module that supports data rates up to 40 Gbps. QSFP+ modules can be used for various network protocols and media types, such as Ethernet, Fibre Channel, InfiniBand, or optical fiber. QSFP+ modules have a 38-pin edge connector and can be inserted into a QSFP+ port on a switch or a server. SFP+ (Small Form-factor Pluggable Plus) is a type of connector that supports data rates up to 10 Gbps, but not 40 Gbps. SFP+ modules have a 20-pin edge connector and can be inserted into an SFP+ port on a switch or a server. GBIC (Gigabit Interface Converter) is an older type of connector that supports data rates up to 1 Gbps, but not 40 Gbps. GBIC modules have an SC duplex connector and can be inserted into a GBIC port on a switch or a server. SFP (Small Form-factor Pluggable) is another older type of connector that supports data rates up to 1 Gbps or 4

Gbps, but not 40 Gbps. SFP modules have an LC duplex connector and can be inserted into an SFP port on a switch or a server. Reference: https://www.howtogeek.com/190014/virtualization-basics-understanding- techniques-and-fundamentals/ https://www.howtogeek.com/428483/what-is-end-to-end- encryption-and-why-does-it-matter/ https://www.howtogeek.com/443611/how-to-encrypt-your- macs-system-drive-removable-devices-and-individual-files/

Question: 143

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

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

Answer: D

Explanation:

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

Question: 144

Users ate experiencing issues when trying to access resources on multiple servers. The servers are virtual and run on an ESX server. A systems administrator is investigating but is unable to connect to any of the virtual servers. When the administrator connects to the host, a purple screen with while letters appears. Which of the following troubleshooting steps should the administrator perform FIRST?

A. Check the power supplies
B. Review the log files.
C. Reinstall the ESX server.
D. Reseat the processors.

Answer: B

Explanation:

A purple screen with white letters on an ESX server indicates a kernel panic, which is a fatal error thatcauses the system to crash and stop functioning3. The first troubleshooting step that an administrator should perform is to review the log files, which may contain information about the cause of the error, such as hardware failures, software bugs, or configuration issues4. Checking the power supplies (A) may not be relevant, as the system is still displaying a screen. Reinstalling the ESX server © or reseating the processors (D) are drastic measures that may result in data loss or further damage, and should only be attempted after ruling out other possible causes.

Reference: 3 https://kb.vmware.com/s/article/1014508 4 https://www.altaro.com/vmware/vmware- esxi-purple-screen-death/

Question: 145

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. Reference: 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.htmlhttps://www.techopedia.com/definition/24195/offsite-storage

Question: 146

A technician is working on a Linux server. The customer has reported that files in the home directory are missing. The /etc/ f stab file has the following entry: nfsserver:/home /home nfs defaults 0 0
However, a df -h /home command returns the following information:
/dev/sda2 10G 1G 9G 10% /home
Which of the following should the technician attempt FIRST to resolve the issue?

A. mkdir /home

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

Answer: B

Explanation:

The /etc/fstab file contains the information about the file systems that are mounted automatically at boot time or on demand. The entry nfsserver:/home /home nfs defaults 0 0 indicates that the /home directory on the local server is mounted from the /home directory on a remote server called nfsserver using the NFS protocol. However, the df -h /home command shows that the /home directory is actually mounted from a local partition /dev/sda2, which may not contain the user's files.

This means that the NFS mount failed or was overridden by another mount. To resolve the issue, the technician should attempt to unmount the local partition using umount nfsserver:/home, which will detach the /home directory from /dev/sda2. Then, the technician should try to mount the NFS share again using mount /home, which will attach the /home directory to nfsserver:/home according to the /etc/fstab entry12. Creating a new directory (A) or removing an existing one © would not help, as they would not affect the mount point. Mounting /home (D) without unmounting it first would not work, asit would result in an error that the mount point is busy3.

Reference: 1 https://askubuntu.com/questions/374870/home-directory-not-being-created 2 https://www.techrepublic.com/article/how-to-properly-automount-a-drive-in-ubuntu- linux/ 3 https://serverfault.com/questions/587855/cannot-find-home-directory-on-linux-server

Question: 147

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

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

Answer: C

Explanation:

An IP KVM (keyboard, video, mouse) is a device that allows remote access and control of multiple servers over a network using a web browser or a clientsoftware. An IP KVM is a cost-effective option to administer and troubleshoot network problems locally on the servers, as it eliminates the need for physical presence or dedicated hardware for each server. A management port (A) is a network interface that is used for out-of-band management of network devices, such as routers or switches. A management port does not provide local access to servers. A crash cart (B) is a mobile unit that contains a monitor, keyboard, mouse, and other tools for troubleshooting servers in a data center. A crash cart requires physical access to each server and may not be cost-effective for many racks of servers. A KVM (D) is a device that allows switching between multiple

servers using a single keyboard, video, and mouse. A KVM does not provide remote access over a network and requires physical connection to each server. Reference: https://www.enterprisestorageforum.com/management/best-data-storage-solutions-and-software- 2021/https://www.microsoft.com/en-us/microsoft-365/business-insights-ideas/resources/cloud- storage-vs-on-premises-servers

Question: 148

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

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

Answer: C

Explanation:

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

Question: 149

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. Reference: 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

Question: 150

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

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

Answer: B,E

Explanation:

These are two measures that should be taken when a data breach occurs. A data breach is an unauthorized or illegal access to confidential or sensitive data by an internal or external actor. A data breach can result in financial losses, reputational damage, legal liabilities, and regulatory penalties for the affected organization. Disclosing the incident is a measure that involves informing the relevant stakeholders, such as customers, employees, partners, regulators, and law enforcement, about the nature, scope, and impact of the data breach. Disclosing the incident can help to mitigate the negative consequences of the data breach, comply with legal obligations, and restore trust and confidence. Identifying the exploited vulnerability is a measure that involves investigating and analyzing the root cause and source of the data breach. Identifying the exploited vulnerability can help to prevent further data loss, remediate the security gaps, and improve the security posture of the organization. Restoring the data from backup is a measure thatinvolves recovering the lost or corrupted data from a secondary storage device or location. However, this does not address the underlying issue of how the data breach occurred or prevent future breaches. Disabling unnecessary ports is a measure that involves closing or blocking network communication endpoints that are not required for legitimate purposes. However, this does not address how the data

breach occurred or what vulnerability was exploited. Running an antivirus scan is a measure that involves detecting and removing malicious software from a system or network. However, this does not address how the data breach occurred or what vulnerability was exploited. Moving the data to a different location is a measure that involves transferring the data to another storage device or location that may be more

secure or less accessible. However, this does not address how the data breach occurred or what vulnerability was exploited. Reference: 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/

Question: 151

DRAG DROP

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

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

INSTRUCTIONS

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

a. PDU selections must be changed using the pencil icon.

b. VM Hosts 1 and 2 and Mail Relay can be moved between racks.

c. Certain devices contain additional details

Data Center Racks 1 and 2

Answer:

Explanation:

Question: 152

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

A. Hardware failure
B. Marware
C. Data corruption
D. Insider threat

Answer: D

Explanation:

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

Question: 153

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

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

Answer: A

Explanation:

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

Reference:

https://www.techopedia.com/definition/16293/mantrap https://www.techopedia.com/definition/1437/bollard https://www.techopedia.com/definition/23961/geofencing

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

Topic 2, Exam Set B

Question: 154

A technician learns users are unable to tog 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 theyare known-working LDAP credentials. The server has access to the LDAP host if it can communicate over the network and is configured correctly. Reference: 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

Question: 155

Which of the following BEST describes a warm site?

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

Answer: D

Explanation:

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

Question: 156

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 duplicatesdata 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. Reference:

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/

Question: 157

The management team at a healthcare organization is concerned about being able to access the dairy vital records if there is an IT disaster that causes both servers and the network lo be ofline. Which of the following backup types can the organization use to mitigate this risk?

A. Tape
B. Cloud
C. Disk
D. Print

Answer: D

Explanation:

A print backup is a type of backup that can be used to mitigate the risk of being unable to access the daily vital records if there is an IT disaster that causes both servers and the network to be ofline. A print backup is a backup that involves printing out the data on paper and storing it in a secure location. A print backup can provide ofline access to the data without relying on any hardware or software components that may be affected by the disaster. However, a print backup has some drawbacks such as high cost, low efficiency, low security, and environmental impact. A tape backup is a type of backup that involves storing the data on magnetic tape cartridges that can be accessed using a tape drive or a tape library. A tape backup can provide ofline access to the data with high capacity, low cost, and long durability, but it requires special equipment and software that may not be available during a disaster. A cloud backup is a type of backup that involves storing the data on remote servers or platforms that can be accessed over the internet using a web browser or an application. A cloud backup can provide online access to the data with high scalability, flexibility, and security, but it requires network connectivity and bandwidth that may not be available during a disaster. A disk backup is a type of backup that involves storing the data on hard disk drives or solid state drives that can be accessed using a computer or a device. A disk backup can provide online or ofline access to the data with high performance, reliability, and portability, but it requires compatible hardware and software that may not be available during a disaster. Reference: https://www.howtogeek.com/428483/what-is-end-to-end-encryption-and-why-does-it-matter/ https://www.howtogeek.com/443611/how-to-encrypt-your-macs-system-drive-removable-devices- and-individual-files/ 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

Question: 158

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

Answer: A

Explanation:

A tabletop exercise is a type of disaster recovery testing exercise that is primarily used to discuss incident response strategies for critical systems without affecting production data. A tabletop exercise is a discussion-based session where team members meet in an informal, classroom setting to review their roles and responsibilities during an emergency and their responses to a hypothetical scenario. A facilitator guides the participants through the discussion and evaluates the strengths and weaknesses of the preparedness program. A tabletop exercise does not involve any actual deployment of resources oractivation of systems12. 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. Reference: 1 https://www.ready.gov/testing- exercises 2 https://www.ready.gov/exercises

Question: 159

A server technician downloaded new firmware from the manufacturer's website. The technician then attempted to install the firmware on the server, but the installation failed, stating the file is potentially corrupt. Which of the following should the technician have checked prior to installing the firmware?

A. DLF configuration
B. MBR failure
C. ECC support
D. MD5 checksum

Answer: D

Explanation:

A MD5 checksum is a value that is calculated from a file using a cryptographic hash function. A MD5 checksum is used to verify the integrity of a file by comparing it with the original value provided by the manufacturer or the source. If the MD5 checksums match, it means that the file is authentic and has not been corrupted or tampered with. If the MD5 checksums do not match, it means that the file is potentially corrupt or malicious and should not be installed12. A DLF configuration (A) is a setting that determines how a dynamic link library (DLL) is loaded into memory and executed by an application. A DLF configuration does not check the integrity of a file. A MBR failure (B) is a problem that occurs when the master boot record (MBR) of a disk is damaged or corrupted, preventing the system from booting. A MBR failure does not check the integrity of a file. ECC support © is a feature that enables error-correcting code (ECC) memory to detect and correct data errors in RAM. ECC support does not check the integrity of a file.Reference: 1 https://www.comparitech.com/net- admin/file-integrity-monitoring-tools/ 2 https://csrc.nist.gov/CSRC/media/Presentations/Firmware- Integrity-Verification-Monitoring-and-Re/images-media/day2_demonstration_330-420.pdf

Question: 160

A technician needs to install a Type 1 hypervisor on a server. The server has SD card slots, a SAS controller, and a SATA controller, and it is attached to a NAS. On which of the following drive types should the technician install the hypervisor?

A. SD card

B. NAS drive
C. SATA drive
D. SAS drive

Answer: A

Explanation:

A SD card is a type of flash memory card that can be used to store data and run applications. A SD card can be used to install a Type 1 hypervisor on a server, as it provides fast boot time, low power consumption, and high reliability. A Type 1 hypervisor runs directly on the underlying computer's physical hardware, interacting directly with its CPU, memory, and physical storage. For this reason, Type 1 hypervisors are also referred to as bare-metal hypervisors. A Type 1 hypervisor takes the place of a host operating system and VM resources are scheduled directly to the hardware by the hypervisor123. A NAS drive (B) is a type of network-attached storage (NAS) device that provides shared access to files and data over a network. A NAS drive cannot be used to install a Type 1 hypervisor on a server, as it requires a network connection and a host operating system to function. A SATA drive © is a type of hard disk drive (HDD) or solid state drive (SSD) that uses the Serial ATA (SATA) interface to connect to a computer. A SATA drive can be used to install a Type 1 hypervisor on a server, but it may have some disadvantages compared to a SD card, such as slower boot time, higher power consumption, and lower reliability. A SAS drive (D) is a type of hard disk drive (HDD) or solid state drive (SSD) that uses the Serial Attached SCSI (SAS) interface to connect to a computer. A SAS drive can also be used to install a Type 1 hypervisor on a server, but it may have similar disadvantages as a SATA drive, and it may also be more expensive and less compatible than a SD card.Reference: 1 https://phoenixnap.com/kb/what-is-hypervisor-type-1-
2 2 https://www.ibm.com/topics/hypervisors 3 https://www.redhat.com/en/topics/virtualization/w hat-is-a-hypervisor

Question: 161

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

A. nbtstat B. netuse C. tracert D. netstst

Answer: B

Explanation:

The net use command is a Windows command that can be used to create a batch script to map multiple shares. The net use command can connect or disconnect a computer from a shared resource, such as a network drive or a printer, or display information about computer connections. The syntax of the net use command is:
net use [devicename | *] [\\computername\sharename[\u0003volume] [password | *]] [/user:[domainname\]username] [/user:[dotted domain name\]username]
[/user:[[username@dotted domain name] [/savecred] [/smartcard] [{/delete | /persistent:{yes | no}}]

where:
devicename = the drive letter or printer port to assign to the shared resource computername = the name of the computer that provides access to the shared resource sharename = the name of the shared resource password = the password needed to access the shared resource /user = specifies a different username to make the connection /savecred = stores the provided credentials for future use /smartcard = uses a smart card for authentication /delete = cancels a network connection and removes the connection from the list of persistent connections /persistent = controls whether the connection is restored at logon
To create a batch script to map multiple shares, you can use the net use command with different drive letters and share names, for example:
net use W: \\computer1\share1 net use X: \\computer2\share2 net use Y: \\computer3\share3 You can also add other options, such as passwords, usernames, or persistence, as needed. To save the batch script, you can use Notepad or any text editor and save the file with a .bat extension12. Reference: 1 https://docs.microsoft.com/en-us/windows-server/administration/windows- commands/net-use 2 https://www.watchingthenet.com/create-a-batch-file-to-map-drives-
folders.html

Question: 162

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

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

Answer: D

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

Question: 163

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.

C. 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. Reference: 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

Question: 164

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. 0

B. 1

C. 5

D. 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 blocklevel striping with one parity block

distributed across all member disks. RAID 5 can tolerate one drive failure without data loss, but not two.

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

Question: 165

A senior administrator instructs a technician to run the following script on a Linux server:

for i in {1..65536); do echo Si; telnet localhost $i; done

The script mostly returns the following message: Connection refused. However, there are several entries in the console display that look like this:

80
Connected to localhost
443
Connected to localhost

Which of the following actions should the technician perform NEXT?

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

Answer: A

Explanation:

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

https://phoenixnap.com/kb/telnet-windows

https://www.techopedia.com/definition/23337/http-port-80 https://www.techopedia.com/definition/23336/https-port-443

Question: 166

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 SANbased shared storage. Fibre Channel is a high-speed network technology that provides reliable and secure data transfer betweenservers 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. Reference: 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

Question: 167

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.Reference:https://docs.microsoft.com/en-us/windows-server/administration/windows- commands/sc-create

Question: 168

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

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

Answer: A

Explanation:

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

Question: 169

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

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

Answer: A

Explanation:

Open-source software is software that is freely available and can be modified and distributed by anyone. It usually requires very little technical support and has no licensing fees. Therefore, it would

be the lowest cost solution for an application that does not need much support.Reference: https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam- objectives (Objective 2.3)

Question: 170

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.Reference: https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam- objectives (Objective 4.1)

Question: 171

An administrator is troubleshooting performance issues on a server that was recently upgraded. The administrator met with users/stakeholders and documented recent changes in an effort to determine whether the server is better or worse since the changes. Which of the following would BEST help answer the server performance question9

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

Answer: B

Explanation:

A server baseline is a set of metrics that represents the normal performance and behavior of a server

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

Question: 172

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

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

Answer: C

Explanation:

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

Question: 173

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.Reference: https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam- objectives (Objective 3.3)

Question: 174

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

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

Answer: C

Explanation:

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

Question: 175

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.Reference: https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-objectives (Objective 3.2)

Question: 176

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.Reference:https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/netstat

Question: 177

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

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

Answer: A

Explanation:

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

Question: 178

An administrator is alerted to a hardware failure in a mission-critical server. The alert states that two drives have failed. The administrator notes the drives are in different RAID 1 arrays, and both are hot- swappable. Which of the following steps will be the MOST efficient?

A. Replace one drive, wait for a rebuild, and replace the next drive.
B. Shut down the server and replace the drives.
C. Replace both failed drives at the same time.
D. Replace all the drives in both degraded arrays.

Answer: C

Explanation:

Since both drives are in different RAID 1 arrays and both are hot-swappable, the most efficient step is to replace both failed drives at the same time. This can minimize the downtime and avoid unnecessary reboots. RAID 1 provides mirroring, which means that data is duplicated on both drives in the array. Therefore, replacing one drive will not affect the data on the other drive or the functionality of the array.Reference:https://en.wikipedia.org/wiki/Standard_RAID_levels#RAID_1

Question: 179

A server administrator made a change in a server's BIOS in an attempt to fix an issue with the OS not turning on. However, the change did not successfully correct the issue. Which of the following should the server administrator do NEXT?

A. Reinstall the server OS in repair mode while maintaining the data.
B. Flash the BIOS with the most recent version.
C. Reverse the latest change made to the server and reboot.
D. Restart the server into safe mode and roll back changes.

Answer: C

Explanation:

The best practice for troubleshooting is to follow a logical and systematic process that involves identifying the problem, establishing a theory of probable cause, testing the theory, establishing a plan of action, implementing the solution, verifying functionality, and documenting findings. Since the problem occurred after a change in the server's BIOS, the most likely cause is that the change was incompatible or incorrect for the OS. Therefore, the next step should be to reverse the latest change made to the server and reboot to see if that fixes the issue.Reference: https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam- objectives (Objective 4.3)

Question: 180

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

Question: 181

A technician is setting up a small office that consists of five Windows 10 computers. The technician has been asked to use a simple IP configuration without manually adding any IP addresses. Which of the following will the technician MOST likely use for the IP address assignment?

A. Static
B. Router-assigned
C. APIPA
D. DHCP

Answer: D

Explanation:

DHCP stands for Dynamic Host Configuration Protocol and it is a network protocol that automatically assigns IP addresses and other network configuration parameters to devices on a network. DHCP can help simplify IP configuration without manually adding any IP addresses. DHCP works by using a DHCP server that maintains a pool of available IP addresses and leases them to devices that request them. The devices can renew or release their IP addresses as needed.Reference: https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-

objectives (Objective 2.1)

Question: 182

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.Reference: https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam- objectives (Objective 2.2)

Question: 183

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.Reference: https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam- objectives (Objective 3.1)

Question: 184

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.

Reference:https://linux.die.net/man/1/chmod

Question: 185

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

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

Which of the following is MOST likely the issue?

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

Answer: A

Explanation:

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

Question: 186

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.Reference: https://www.comptia.org/training/resources/exam- objectives/comptia-server-sk0-005-exam-objectives (Objective 2.4)

Question: 187

An administrator is installing a new file server that has four drive bays available. Which of the following RAID types would provide the MOST storage as well as disk redundancy?

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

Answer: C

Explanation:

RAID 5 is a RAID level that provides striping with parity, which means that data is distributed across all disks with one disk storing parity information for error correction. RAID 5 can tolerate one disk failure without losing data. RAID 5 provides the most storage as well as disk redundancy out of the

four RAID levels given, since it only uses one disk for parity and the rest for data. For example, if four 200GB drives are used in a RAID 5 array, the total storage capacity would be 600GB (200GB x 3), while in RAID 0 it would be 800GB (200GB x 4), in RAID 1 it would be 200GB (200GB x 1), and in RAID 10 it would be 400GB (200GB x 2).Reference:https://en.wikipedia.org/wiki/Standard_RAID_levels#RAID_5

Question: 188

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.Reference:https://en.wikipedia.org/wiki/Standard_RAID_levels#RAID_5

Question: 189

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.Reference: https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-objectives (Objective 3.3)

Question: 190

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

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

Answer: C

Explanation:

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

Question: 191

A server shut down after an extended power outage. When power was restored, the system failed to start. A few seconds into booting, the Num Lock, Scroll Lock, and Caps Lock LEDs flashed several times, and the system stopped. Which of the following is the MOST likely cause of the issue?

A. The keyboard is defective and needs to be replaced.
B. The system failed before the display card initialized.
C. The power supply is faulty and is shutting down the system.
D. The NIC has failed, and the system cannot make a network connection.

Answer: B

Explanation:

This is the most likely cause of the issue because the keyboard LED flash indicates a POST error code. If the display card is not initialized, the system cannot show any error messages on the screen and will stop booting.
Reference:https://www.computerhope.com/beep.htm#04

Question: 192

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

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

Answer: A

Explanation:

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

Question: 193

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

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

Answer: B

Explanation:

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

Question: 194

A systems administrator is trying to determine why users in the human resources department cannot access an application

server. The systems administrator reviews the application logs but does not see any attempts by the users to access the application. Which of the following is preventing the users from accessing the application server?

A. NAT
B. ICMP
C. VLAN
D. NIDS

Answer: C

Explanation:

This is the most likely cause of preventing the users from accessing the application server because a VLAN is a logical segmentation of a network that isolates traffic based on certain criteria. If the human resources department and the application server are on different VLANs, they will not be able to communicate with each other unless there is a router or a switch that can route between VLANs.Reference:https://www.cisco.com/c/en/us/support/docs/lan-switching/inter-vlan-routing/41860-howto-L3-intervlanrouting.html

Question: 195

An administrator is only able to log on to a server with a local account. The server has been successfully joined to the domain and can ping other servers by IP address. Which of the following locally defined settings is MOST likely misconfigured?

A. DHCP
B. WINS
C. DNS
D. TCP

Answer: C

Explanation:

This is the most likely misconfigured setting because DNS is the service that resolves hostnames to IP addresses and vice versa. If the DNS server is incorrect or unreachable, the administrator will not be able to log on to the server with a domain account because the server will not be able to authenticate with the domain controller.Reference:https://docs.microsoft.com/en- us/troubleshoot/windows-server/networking/dns-troubleshooting

Question: 196

An administrator is investigating several unexpected documents and video files that recently appeared in a network share.

The administrator checks the properties of the files and sees the author's name on the documents is not a company employee. The administrator questions the other users, but no one knows anything about the files. The administrator then checks the log files and discovers the FTP protocol was used to copy the files to the server. Which of the following needs to be done to prevent this from happening again?

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

Answer: D

Explanation:

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

Question: 197

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

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

Answer: D

Explanation:

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

Question: 198

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

A. Private

B. Hybrid
C. Community
D. Public

Answer: B

Explanation:

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

Question: 199

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:

This is the best description of a synthetic full backup method because it creates a full backup by combining previous incremental backups with the latest backup. An incremental backup copies only the files that have changed since the last backup, while a full backup copies all the files. A synthetic full backup reduces the storage space and network bandwidth required for backups, while also simplifying the restore process by using a single file.Reference:https://www.veritas.com/support/en_US/doc/129705091-129705095-0/br731_wxrt- tot_v131910378-129705095

Question: 200

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

A. Unplug the power cables from the redundant power supplies, leaving just the minimum required.

B. Convert the physical servers to the hypervisors and retire the ten servers.
C. Reimage the physical servers and retire all ten servers after the migration is complete.
D. Convert the ten servers to power-efficient core editions.

Answer: B

Explanation:

This option would reduce power consumption and maintenance costs by consolidating the physical servers into virtual machines on the hypervisors. Thiswould also free up space and resources in the

data center. The other options would either not address the request, increase power consumption, or require more maintenance.

Question: 201

A server administrator encounters some issues with the server OS after applying monthly patches.
Which of the following troubleshooting steps should the administrator perform?

A. Implement rollback procedures.
B. Upgrade the drivers.
C. Reinstall the OS.
D. Reboot the server.

Answer: A

Explanation:

This option would restore the server OS to a previous state before applying the monthly patches. This would help troubleshoot the issues caused by the patches and determine if they are compatible with the server OS. The other options would either not address the issues, cause data loss, or require more time and resources

Question: 202

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

A. Telnet connections are disabled on the server.
B. Role-based access control is misconfigured.
C. There are misconfigured firewall rules.

D. Group policies have not been applied.

Answer: C

Explanation:

This is the most likely cause of the issue because firewall rules can block or allow traffic based on source, destination, port, protocol, or other criteria. If the firewall rules are not configured properly, they can prevent remote users from accessing the cloud application server, while allowing on-site

users to access it.Reference:https://docs.microsoft.com/en-us/azure/virtual-network/security- overview

Question: 203

A technician is monitoring a server and notices there is only one NIC plugged in. but the server has two. The NIC is oversaturated, and the technician would like to increase the available bandwidth. Which of the following solutions would be the BEST option to increase the speed of this NIC?

A. Link aggregation
B. Heartbeat
C. Most recently used
D. Active-active

Answer: A

Explanation:

This is the best solution to increase the speed of the NIC because link aggregation is a technique that combines multiple physical network interfaces into a single logical interface. This can increase the bandwidth, redundancy, and load balancing of network traffic. Link aggregation requires both the server and the switch to support it and be configured accordingly.Reference:https://www.cisco.com/c/en/us/support/docs/lan- switching/etherchannel/12023-4.html

Question: 204

A user logs in to a Linux server and attempts to run the following command:

sudo emacs /root/file

However the user gets the following message:

User userid is not allowed to execute Temacs' on this server. Which of the following would BEST allow the user to find out which commands can be used?

A. visudo I grep userid
B. sudo -1 -U userid

C. cat /etc/passwd
D. userlist I grep userid

Answer: B

Explanation:

This is the best command to find out which commands can be used by a user with sudo privileges because it lists the allowed and forbidden commands for a given user or role. The -l option stands for list, and the -U option specifies the user name. The output of this command will show what commands can be executed with sudo by that user on that server.Reference:https://www.sudo.ws/man/1.8.13/sudo.man.html

Question: 205

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

A. Use an MFAout-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.Reference:https://docs.microsoft.com/en-us/azure/active- directory/authentication/concept-mfa-howitworks

Question: 206

Which of the following refers to the requirements that dictate when to delete data backups?

A. Retention policies.
B. Cloud security impact
C. Off-site storage
D. Life-cycle management

Answer: A

Explanation:

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

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

Question: 207

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: B,C

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

Question: 208

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.

Question: 209

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

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

Answer: B

Explanation:

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

Question: 210

A technician re working on a Linux server and re trying to access another server over the network. The technician gets server notfoundmessage when trying to execute ping servername but no error messages when usingpingservername. Domain.com. Which ofthefollowing should the technician do to resolve the error?

A. Configure the domain search variable
B. Change the permissions on resolv. conf

C. Configure the DNS address D. Modify nsswitch. Conf.

Answer: A

Explanation:

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

Question: 211

Joe. a user m the IT department cannot save changes to a sensitive file on a Linux server. An 1s -1& snows the following listing;

Which of the following commands would BEST enable the server technician to allow Joe tohaveaccess 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 Reference: [Linux chmod command]

Question: 212

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

implement with no downtime?

A. Configure a RAID array.
B. Replace the current drives with higher-capacity disks.

C. Implement FCoE for more storage capacity.
D. Connect the server to a SAN

Answer: D

Explanation:

A SAN (Storage Area Network) is a network of storage devices that can provide shared storage capacity to multiple servers. By connecting the server to a SAN, theadministrator 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 Reference: [What is a SAN and how does it differ from NAS?]

Question: 213

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

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

Answer: D

Explanation:

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

Question: 214

A server administrator recently installed a kernel update to test functionality Upon reboot, the administrator determined the new kernel was not compatible with certain server hardware and was unable to uninstall the update. Which of the following should the administrator do to mitigate further issues with the newly instated kernel version?

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

Answer: A

Explanation:

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

Question: 215

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 Reference: [Incremental vs Differential Backup]

Question: 216

Users report they ate unable to access an application after a recent third-party patch update. The physical server that is hosting the application keeps crashing on reboot. Although the update was installed directly from the manufacturer's support website as recommended it has now been recalled and removed from the website as the update unintentionally installed unauthorized software after a reboot. Which of the following steps should the administrator perform to restore access to the application while minimizing downtime? (Select TWO)

A. Uninstall recent updates.
B. Reimage the server with a different OS.
C. Run a port scan to verify open ports.
D. Enable a GPO to uninstall the update.
E. Scan and remove any malware.
F. Reformat the server and restore the image from the latest backup.

Answer: E,F

Explanation:

The most likely cause of the server crashing and the application being inaccessible is that the unauthorized software installed by the update is malware that corrupted the system files or compromised the security of the server. To restore access to the application while minimizing downtime, the administrator should scan and remove any malware from the server, and then reformat the server and restore the image from the latest back-up. This will ensure that the server is clean and has a working configuration of the application. Verified Reference: [How to Remove Malware from a Server]

Question: 217

Which of the following backup types should be chosen for database servers?

A. Differential
B. Incremental
C. Synthetic full
D. Open file

Answer: C

Explanation:

A synthetic full backup is a type of backup that combines a full backup with one or more incremental backups to create a new full backup without accessing the source data. This type of backup is suitable for database servers, as it reduces the backup window, minimizes the impact on the server performance, and provides faster recovery time. Verified Reference:

[Synthetic Full Backup]

Question: 218

An administrator is troubleshooting a failure in the data center in which a server shut down/turned off when utility power was lost The server had redundant power supplies. Which of the following is the MOST likely cause of this failure?

A. The UPS batteries were overcharged.
B. Redundant power supplies require 220V power
C. Both power supplies were connected to the same power feed
D. The power supplies werenot cross-connected

Answer: C

Explanation:

The most likely cause of this failure is that both power supplies were connected to the same power feed, which means that they both lost power when utility power was lost. To prevent this from happening, redundant power supplies should be connected to different power feeds, preferably from different sources, such as a UPS or a generator. Verified Reference: [Redundant Power Supply Best Practices]

Question: 219

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

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

Answer: C

Explanation:

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

Question: 220

While running a local network security scan an administrator discovers communication between clients and one of the web servers is happening in cleartext. Company policy requires all communication lo be encrypted. Which of the following ports should be closed to stop the deartext communication?

A. 21
B. 22
C. 443
D. 3389

Answer: A

Explanation:

Port 21 is used for FTP (File Transfer Protocol), which is a protocol that transfers files between servers and clients in cleartext, meaning that anyone can intercept and read the data. To stop this communication, port 21 should be closed on the web server and replaced with a secure protocol, such as SFTP (Secure File Transfer Protocol) or FTPS (File Transfer Protocol Secure), which use encryption to protect the data. Verified Reference: [FTP vs SFTP vs FTPS]

Question: 221

The accounting department needs more storage and wants to retain the current data for quick readwrite access. The accounting server does not have any internet drive bays available to keep both disks however the server does have USB 3 0 and eSATA ports available. Which of the following is the BEST way to accomplish the department's goals?

A. Copy the existing data to an external USB 3 0 enclosure.
B. Place the existing data on a DVD and use the internal DVD-ROM drive.
C. Transfer the existing data to an external eSATA enclosure.
D. Move the existing data to a new. larger internal hard drive.

Answer: C

Explanation:

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

Question: 222

Which ofthefollowingshouldatechnician verify FIRST before decommissioning and wipingafile server?

A. Themedia destructionmethod
B. Therecycling poke?
C. Asset management documentation
D. Non-utilization

Answer: D

Explanation:

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

Question: 223

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

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

Answer: B

Explanation:

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

Question: 224

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

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

Answer: C

Explanation:

RTO (Recovery Time Objective) is a metric that measures how much downtime an organization can tolerate during an unplanned outage before it affects its business continuity and reputation. RTO is usually expressed in hours or minutes and is determined by the criticality of the business processes and the impact of the outage on the revenue, customers, and stakeholders. RTO helps to define the recovery strategy and the resources needed to restore the normal operations as quickly as possible. Verified Reference: [RTO vs RPO]

Question: 225

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

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

Answer: C

Explanation:

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

Question: 226

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

A. Network encapsulation

B. Off-site data

C. Secure FTP

D. Data in transit

Answer: D

Explanation:

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

Question: 227

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

A. Batch

B. VBS

C. Bash

D. PowerShell

Answer: D

Explanation:

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

Question: 228

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

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

Answer: B

Explanation:

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

Question: 229

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

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

Answer: D

Explanation:

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

Question: 230

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

A. Installing blanking panels in all the empty rack spaces

B. installing an additional POU and spreading out the power cables
C. Installing servers on the shelves instead of sliding rails
D. installing front bezels on all the server's m the rack

Answer: A

Explanation:

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

Question: 231

Which of the following asset management documents is used to identify the location of a serves within a data center?

A. Infrastructure diagram
B. Workflow diagram
C. Rack layout
D. Service manual

Answer: C

Explanation:

A rack layout is a document that shows the physical location and arrangement of servers and other devices within a rack. It can include information such as server names, IP addresses, power consumption, and cable connections. A rack layout can help identify and locate servers easily and efficiently in a data center. Verified Reference: [Rack layout], [Data center]

Question: 232

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 Reference: [PowerShell security commands], [Security policy]

Question: 233

A developer is creating a web application that will contain five web nodes. The developer's main goal is to ensure the application is always available to the end users. Which ofthefollowing should the developer use when designing the web application?

A. Round robin
B. Link aggregation
C. Network address translation
D. Bridged networking

Answer: A

Explanation:

Round robin is a load balancing technique that distributes requests among multiple web nodes in a circular order. It ensures that each web node receives an equal amount of requests and improves the availability and performance of the web application. Verified Reference: [Round robin], [Load balancing]

Question: 234

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

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

Answer: D

Explanation:

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

Question: 235

Which of thefollowinglicensing concepts is based on the number of logical processorsaserverhas?

A. Per core
B. Per socket
C. Per instance
D. Per server

Answer: A

Explanation:

Per core licensing is based on the number of logical processors a server has. A logical processor is either a physical core or a virtual core created by hyperthreading. Per core licensing requires purchasing a license for each logical processor on the server. Verified Reference: [Per core licensing], [Logical processor]

Question: 236

An application server's power cord was accidentally unplugged. After plugging the cord back in the server administrator notices some transactions were not written to the disk array. Which of the following is the MOST likely cause of the issue?

A. Backplane failure
B. CMOS failure
C. Misconfigured RAID
D. Cache battery failure

Answer: D

Explanation:

A cache battery is a battery that provides backup power to the cache memory of a disk array controller. The cache memory stores data that is waiting to be written to the disk array. If the cache battery fails, the data in the cache memory may be lost or corrupted when the power is interrupted. Verified Reference: [Cache battery], [Disk array controller]

Question: 237

A server administrator has received calls regarding latency and performance issues with a file server. After reviewing all logs and server features the administrator discovers the server came with four Ethernet ports, out only one port is currently in use. Which of the following features will enable the use of all available ports using a single IP address?

A. Network address translation
B. in-band management
C. Round robin
D. NIC teaming

Answer: D

Explanation:

NIC teaming is a feature that allows the use of multiple network interface cards (NICs) as a single logical interface with a single IP address. It can improve the network performance, bandwidth, and redundancy of a server. Verified Reference: [NIC teaming], [Network interface card]

Question: 238

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

A. increase theaudit. log file size in the appropriate confutation file.
B. Decrease the duration of the log rotate cycle tor theaudit. log file.
C. Remove the tog rotate directive from the audit .log We configuration.
D. Move the audit. leg files to a remote syslog server.

Answer: A

Explanation:

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

Question: 239

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

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

Answer: C

Explanation:

DLP (Data Loss Prevention) is a security measure that prevents unauthorized copying, transferring, or leaking of sensitive data from a server or a network. It can block or alert the user when they try to

copy a file to certain locations, such as a network share or a USB flash drive, based on predefined policies and rules.
Verified Reference: [DLP], [Data loss]

Question: 240

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

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

Answer: B

Explanation:

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

Question: 241

A technician recently applied a critical OS patch to a working sever. After rebooting, the technician notices the server Is unable to connect to a nearby database server. The technician validates a connection can be made to thedatabasefrom another host. Which of the following is the best NEXT step to restore connectivity?

A. Enable HIDS.
B. Change the service account permissions.
C. Check the host firewall I rule.
D. Roll back the applied patch.

Answer: C

Explanation:

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

Question: 242

A server administrator is instating a new server in a data center. The administrator connects the server to a midplane but does not connect any cables Which of the following types of servers is the administrator MOST likely installing?

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

Answer: D

Explanation:

A blade server is a type of server that is installed in a chassis or an enclosure that provides power, cooling, networking, and management features. The blade server does not have any cables attached to it, as it connects to the chassis through a midplane or a backplane. A blade server can save space, energy, and cost compared to other types of servers. Verified Reference: [Blade server], [Chassis]

Question: 243

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

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

Answer: D

Explanation:

GPT (GUID Partition Table) is a partitioning scheme that allows creating partitions on large hard drives (more than 2 TB). It supports up to 128 partitions per drive and uses 64-bit addresses to locate them. MBR (Master Boot Record) is an older partitioning scheme that has limitations on the size and number of partitions (up to 4 primary partitions or 3 primary and 1 extended partition per drive). To provision a new 4 TB drive, the technician should partition it using GPT. Verified Reference: [GPT], [MBR]

Question: 244

A systems administrator has several different types of hard drives. The administrator is setting up a MAS that will allow end users to see all the drives within the NAS. Which of the following storage types should the administrator use?

A. RAID array
B. Serial Attached SCSI
C. Solid-state drive
D. Just a bunch of disks

Answer: D

Explanation:

JBOD (Just a Bunch Of Disks) is a storage configuration that combines different types and sizes of hard drives into one logical unit without any RAID level or redundancy. It allows users to see all the drives within the unit as one large storage space.JBOD can utilize all the available capacity of the drives but does not provide any performance or fault tolerance benefits. Verified Reference: [JBOD], [RAID]

Question: 245

Two developers are working together on a project, and they have built out a set of snared 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 costefficiency for users who need access to various applications and data without investing in their own hardware or software. Verified Reference: [Public cloud], [Cloud model]

Question: 246

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

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

Answer: B

Explanation:

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

Question: 247

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

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

D. Connect the virtual administration console.

Answer: C

Explanation:

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

Question: 248

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

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

Answer: B,D

Explanation:

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

Question: 249

A technician has moved a data drive from a new Windows server to an orderWindows 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. Thedisk uses GPT.
B. Thepartition is formatted with ext4.
C. The -partition is formatted with FAT32.
D. Thedisk uses MBn.

Answer: A

Explanation:

GPT (GUID Partition Table) is a partitioning scheme that allows creating partitions on large hard drives (more than 2 TB). It supports up to 128 partitions per drive and uses 64-bit addresses to locate them. However, GPT is not compatible with older versions of Windows, such as Windows XP or Windows Server 2003, which use MBR (Master Boot Record) as the partitioning scheme. If a disk uses GPT, it may not be recognized or accessible by an older Windows server. Verified Reference: [GPT], [MBR]

Question: 250

Which of the following describes a configuration inwinch 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 Reference: [Active-active], [Redundant system]

Question: 251

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 blockingthe ICMP echo reply.
B. The DHCP scope option is incorrect
C. The DNS entriesforthe 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 Reference: [Hosts file], [DNS]

Question: 252

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

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

Answer: D

Explanation:

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

Question: 253

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

A. Biometrics
B. Push notifications
C. Smart carts
D. Physical tokens

Answer: B

Explanation:

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

Question: 254

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

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

Answer: D

Explanation:

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

Question: 255

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

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

Answer: A

Explanation:

A DHCP reservation is a way to assign a specific IP address to a device based on its MAC address, which is a unique identifier for each network interface card (NIC). When the administrator replaced the NIC, the MAC address of the server changed, and the DHCP server no longer recognized it as the same device. Therefore, the DHCP server assigned a different IP address to the server, which was incorrect for the application. To fix this problem, the administrator needs to modify the DHCP reservation to use the new MAC address of the NIC, so that the server can get the proper IP address. A WWNN (World Wide Node Name) is a unique identifier for a Fibre Channel node, which is a device that can communicate over a Fibre Channel network. A WWNN is not related to DHCP or IP addresses, and it is not used for DHCP reservations. Therefore, options B and D are incorrect. Updating the local hosts file with the correct IP address (option C) is also incorrect, because it does not solve the problem of getting the correct IP address from the DHCP server. The hosts file is a local file that maps hostnames to IP addresses, and it is used to override DNS queries. However, it does not affect how the DHCP server assigns IP addresses to devices. Moreover, updating the hosts file manually on every device that needs to communicate with the server is not a scalable or efficient solution.

Reference:

How to reserve IP Address in DHCP Server - Ask Ubuntu

Static IP vs DHCP Reservation - The Tech Journal

How to Configure DHCP Server Reservation in Windows ... - ITIngredients

Question: 256

An administrator is deploying a new secure web server. The only administration method that is permitted is to connect

via RDP. Which of the following ports should be allowed? (Select TWO).

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

Answer: D,F

Explanation:

Port 443 is the default port for HTTPS, which is the protocol used for secure web communication. HTTPS uses SSL/TLS certificates to encrypt the data between the web server and the browser. Port 443 is commonly used for web servers that need to provide secure services, such as online banking, e-commerce, or email. By allowing port 443, the administrator can access the web server's interface and manage its settings1.

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

Question: 257

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

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

Answer: D

Explanation:

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

RAID 1 provides redundancy and fault tolerance by mirroring the data from one disk to another disk.
RAID 1 offers high read performance and data security, but it has low capacity and write performance. RAID 1 requires a minimum of two disks to operate, and it can only tolerate the failure of one disk in the array. If more than one disk fails, all the data in the array is lost2.
RAID 5 provides redundancy and fault tolerance by using one parity block to store information that can be used to reconstruct the data in case of a disk failure. RAID 5 requires a minimum of three disks to operate, and it 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.

Question: 258

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.

Question: 259

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.

Question: 260

A server administrator is implementing an authentication policy that will require users to use a token during login. Which of the following types of authentication is the administrator implementing?

A. Something you are
B. Something you know
C. Something you have
D. Something you do

Answer: C

Explanation:

Something you have is one of the types of authentication methods that relies on a physical object or device that the user possesses to verify their identity. A token is an example of something you have, as it is a small device that generates a one-time password or code that the user enters during login. A token can be a hardware device, such as a key fob or a smart card, or a software application, such as an app on a smartphone or a browser extension. A token provides an additional layer ofsecurity to theauthentication process, as it prevents unauthorized access even if the user's username and password are compromised1.

Question: 261

Corporate policy mandates that logs from all servers be available for review regardless of the state of the server. Which of the following must be configured to comply with this policy?

A. Aggregation
B. Subscription
C. Merging
D. Collection

Answer: A

Explanation:

Aggregation is the process of collecting, standardizing, and consolidating log data from multiple sources into a central location. Aggregation makes it easier to search, analyze, and report on log
data, as well as to comply with security policies and regulations. By aggregating logs from all servers, regardless of their state, the corporate policy can ensure that no log data is lost or inaccessible in case of a server failure or outage

Question: 262

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

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

Answer: C

Explanation:

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

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

Question: 263

The Chief Information Officer of a data center is concerned that transmissions from the building can be detected from the outside. Which of the following would resolve this concern? (Select TWO).

A. RFID
B. Proximity readers
C. Signal blocking
D. Camouflage
E. Reflective glass
F. Bollards

Answer: C,D

Explanation:

Signal blocking is a technique that prevents or reduces the transmission of electromagnetic signals from a building to the outside. Signal blocking can be achieved by using materials that absorb, reflect, or scatter the signals, such as metal, concrete, or mesh. Signal blocking can protect the data center from eavesdropping, interference, or jamming by unauthorized parties1.

Camouflage is a technique that disguises or conceals the appearance of a building to make it less noticeable or identifiable from the outside. Camouflage can be achieved by using colors, patterns, shapes, or vegetation that blend in with the surrounding environment. Camouflage can protect the data center from detection, reconnaissance, or targeting by hostile parties

Question: 264

SIMULATION

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

INSTRUCTIONS

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

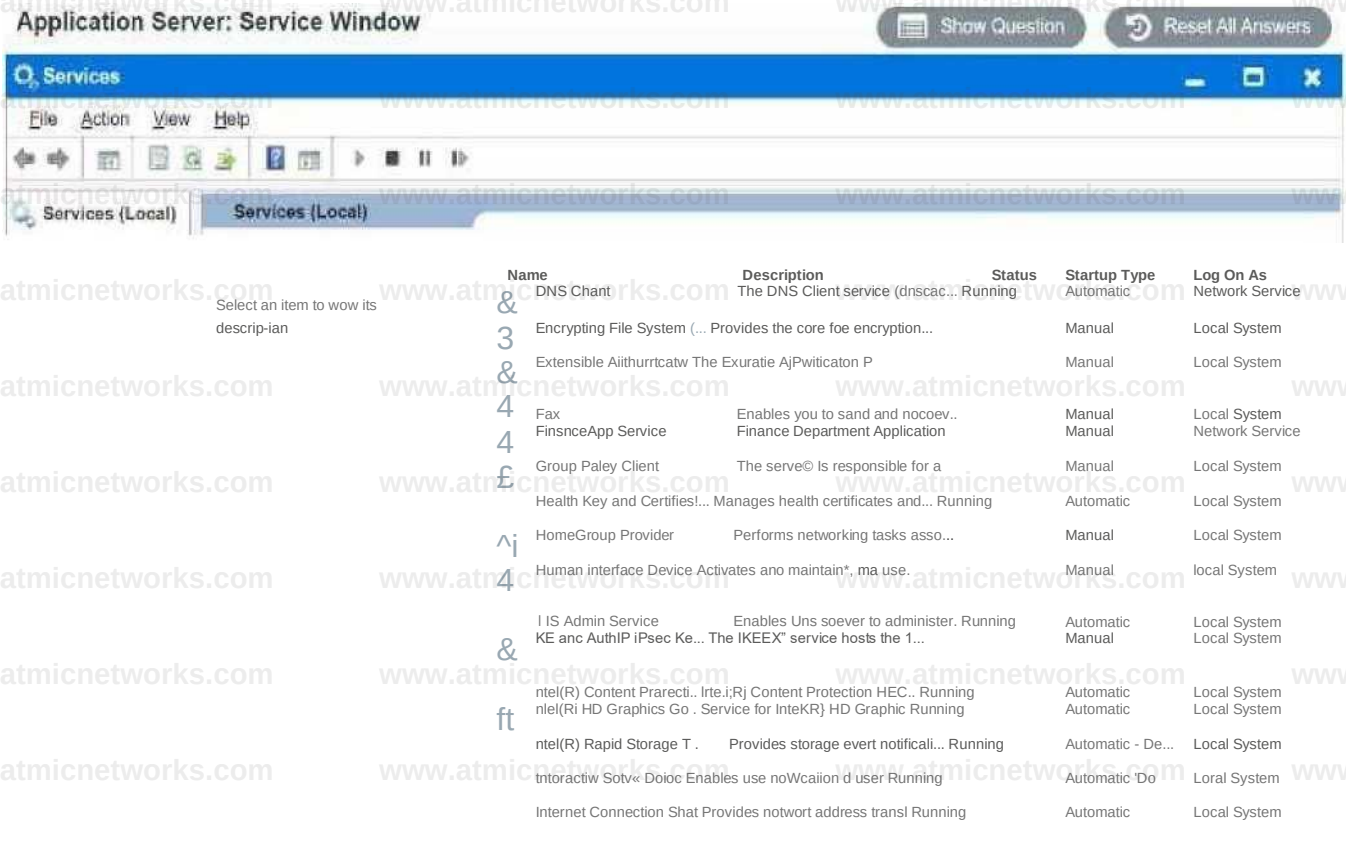

Answer: See the solution in explanation.

Explanation:

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

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

First, you need to start the service, so that the software can run. To do this, you can right-click on the service name and select “Start” from the menu. Alternatively, you can select the service name and click on the “Start” button on the toolbar. You should see a message saying that the service has started successfully.

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

you can select “Automatic” from the drop-down list. Then, click on “OK” to save your changes. By doing these two steps,

you should be able to use the FinanceApp software without any problems. The MarketApp software is running as a service named “MarketApp Service”. The service description says “Provides market data and analysis for the MarketApp software”. The service status is “Running”, which means that the service is running and the software is functional. However, as you reported, the software may crash sometimes, which can cause significant downtime. The service startup type is “Automatic”, which means that the service starts automatically every time the server boots up. The service log on as is “TestAccount”, which is a test account that was probably used for development or testing purposes.

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

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

Question: 265

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

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

Answer: D

Explanation:

The correct answer is D. per-core.

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

power1

Question: 266

The management team has mandated the use of data-at-rest encryption on all corporate servers. Using this encryption paradigm will ensure:

A. website traffic is protected while traversing the internet.
B. files stored on the server are protected against physical theft.
C. attachments that are emailed from this server cannot be intercepted.
D. databases in use are protected from remote hackers.

Answer: B

Explanation:

Data-at-rest encryption is a method of encrypting data while it is stored on a storage device, such as a hard drive, an SSD, or a tape library. This ensures that if the data is stolen or lost, it will be unreadable without the encryption key. Data-at-rest encryption does not protect data while it is in transit over the network, in use by the CPU or memory, or attached to an email.

Question: 267

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

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

Answer: D

Explanation:

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

Question: 268

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 tools12.

Question: 269

A company is implementing a check-in desk to heighten physical security. Which of the following access controls would be

the most appropriate to facilitate this implementation?

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

Answer: D

Explanation:

An access control vestibule, or mantrap, is a type of physical access control that provides a space between two sets of interlocking doors. It is designed to prevent unauthorized individuals from following authorized individuals into facilities with controlled access, such as a check-in desk. The vestibule can be configured to limitthe number of individuals who enter the controlled area and to verify their authorization for physical access1. The other options are incorrect because they are not as effective as an access control vestibule in facilitating the implementation of a check-in desk. Security guards, security cameras, and bollards are useful for monitoring, deterring, or preventing unauthorized access, but they do not provide the same level of control and verification as an access control vestibule

Question: 270

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

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

Answer: C

Explanation:

MRU, or Most Recently Used, is a concept that refers to prioritizing a connection that had previously worked successfully. It is often used in load balancing algorithms to distribute the workload among multiple servers or paths. MRU assumes that the most recently used connection is the most likely to be available and efficient, and therefore assigns the next request to that connection. This can help reduce latency and improve performance12. The other options are incorrect because they do not refer to prioritizing a previous connection. Round robin is a concept that refers to distributing the workload equally

among all available connections in a circular order12. SCP, or Secure Copy Protocol, is a concept that refers to transferring files securely between hosts using encryption3. Link aggregation is a concept that refers to combining multiple physical links into a single logical link to increase bandwidth and redundancy4.

Question: 271

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

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

Answer: C

Explanation:

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

Question: 272

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.

Question: 273

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 CompTIAServer+ (SK0-004) | Pluralsight 2: Introducing the CompTIA Infrastructure Career Pathway

Question: 274

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

Name	Size
01012020.bak	WOMB
01022020.bak	WMB
01032020.bak	5MB
01042020.bak	7MB
01052020.bak	WOMB
01062020.bak	8MB
01072020.bak	WMB

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.

Question: 275

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.

Question: 276

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 addressesin the range of 192.168.0.0 to 192.168.255.2552. To fix this issue, the systems administrator needs tocheck and correct the subnet mask of the new virtual server according to the network configuration.

Question: 277

A server administrator is racking new servers in a cabinet with multiple connections from the servers to power supplies and the network. Which of the following should the administrator recommend to the organization to best address this situation?

A. Rack balancing
B. Cable management
C. Blade enclosure
D. Rail kits

Answer: B

Explanation:

Cable management is the process of organizing, securing, and labeling cables in a server rack or cabinet. Cable management can help improve airflow and cooling, reduce clutter and confusion, prevent damage and interference, and enhance safety and aesthetics123. Cable management can be achieved by using various tools and accessories, such as cable trays, ties, hooks, clips, labels, ducts, and organizers12.

Question: 278

An administrator discovers a misconfiguration that impacts all servers but can be easily corrected. The administrator has a list of affected servers and a script to correct the issue. Which of the following scripting principles should the administrator use to cycle through the list of servers to deliver the needed change?

A. Linked list
B. String
C. Loop
D. Constant

Answer: C

Explanation:

A loop is a programming construct that allows a block of code to be executed repeatedly until a certain condition is met1. A loop can be used to cycle through a list of servers and run a script on each one of them. For example, in Python, a loop can be written as: Python

This code is AI-generated. Review and use carefully. Visit our FAQ for more information. Copy

```
#  Assume servers is a list of server names
forserverinservers:
#  Run the script on the server run_script(server)
```

A loop can help automate the task of correcting the misconfiguration on all servers, saving time and effort.

Question: 279

A storage engineer responds to an alarm on a storage array and finds the battery on the RAID controller needs to be replaced. However, the replacement part will not be available for 14 days. The engineer needs to identify the impact of the failed battery on the system. Which of the following best describes the impact?

A. The read and write performance will be impacted.
B. The read performance will be impacted.
C. The performance will not be impacted.
D. The write performance will be impacted.

Answer: D

Explanation:

A RAID controller battery is used to protect the data in the cache memory of the controller in case of a power failure. The cache memory allows the controller to improve the write performance by buffering the data and writing it to the disk in an optimized way. However, if the battery fails, the controller will switch to write-through mode, which means it will write the data directly to the disk without caching. This will reduce the write performance and increase the latency of the system

Question: 280

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

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

Answer: B

Explanation:

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

Question: 281

An administrator gave Ann modify permissions to a shared folder called DATA, which is located on the company server. Other users need read access to the files in this folder. The current configuration is as follows:

The administrator has determined Ann cannot write anything to the DATA folder using the network. Which of the following would be the best practice to set up Ann's permissions correctly, exposing only the minimum rights required?

A. Option A B. Option B C. Option C D. Option D

Answer: D

Explanation:

Option D is the best practice to set up Ann's permissions correctly, exposing only the minimum rights required. Option D shows that the share permissions on the DATA folder grant Ann Change access, which allows her to read, write, and delete files in the shared folder. The file permissions grant Ann Modify access, which allows her to read, write, execute, and delete files in the folder. This combination of permissions gives Ann the ability to write anything to the DATA folder using the network, as well as to modify and delete existing files. This meets the requirement of giving Ann modify permissions to the shared folder.

Question: 282

Which of the following life-cycle management phases deals with a server that is no longer in operation?

A. End-of-life
B. Disposal
C. Usage
D. Procurement

Answer: A

Explanation:

End-of-life is the phase of lifecycle management that deals with a server that is no longer in operation. End-of-life means that the server has reached the end of its useful life and is no longer supported by the manufacturer or the service provider. End-of-life may also imply that the serveris obsolete, incompatible, or inefficient for the current needs and standards1. End-of-life servers may be decommissioned, recycled, donated, or disposed of according to the organizational policies and environmental regulations

Question: 283

A systems administrator is attempting to install a package on a server. After downloading the package from the internet and trying to launch it, the installation is blocked by the antivirus on the server. Which of the following must be completed before launching the installation package again?

A. Creating an exclusion to the antivirus for the application
B. Disabling real-time scanning by the antivirus
C. Validating the checksum for the downloaded installation package
D. Checking for corruption of the downloaded installation package

Answer: C

Explanation:

A checksum is a value that is calculated from a data set to verify its integrity and authenticity. A checksum can be used to compare a downloaded installation package with the original source to ensure that the package has not been corrupted or tampered with during the download or transmission process. If the checksums match, then the package is safe to install. If the checksums do not match, then the package may be infected with malware or contain errors that could cause installation problems. Therefore, validating the checksum for the downloaded installation package is a necessary step before launching the installation again12

1: CompTIA Server+ Certification Exam Objectives 2: How to Verify File Integrity Using Checksums on Linux

Question: 284

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

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

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

E. 53

Answer: E

Explanation:

Port 53 is the standard port for DNS (Domain Name System) queries and responses. DNS is a service that translates domain names (such as www.example.com) into IP addresses (such as 192.0.2.1) and vice versa. DNS is essential for internet connectivity, as it allows users and applications to access websites and other online resources by using human-readable names instead ofnumerical addresses1.

The DNSSrv server is a DNS server that provides name resolution for the corporate network. If port 53 is blocked on this server, it will not be able to communicate with other DNS servers or clients, and the name resolution will fail. This will prevent users from accessing any websites or online services that rely on domain names, such as web browsers, email clients, or cloud applications. Therefore, port 53 must remain open on DNSSrv to allow DNS traffic to flow.

Question: 285

A technician needs maximum power redundancy while configuring a new server rack. Which of the following should be specified? (Select two).

A. Separate circuits B. Rack balancing C. Blanking panels D. High-voltage PDUs E. KVM placement F. Multiple UPSs

Answer: A,F

Explanation:

Separate circuits and multiple UPSs can provide fail-safe power redundancy for a server rack. Separate circuits mean that the server rack is connected to two or more independent power sources, such as different utility lines or generators. This can prevent a single point of failure in the power supply, as if one circuit goes down, the other circuit can still provide power to the server rack1. Multiple UPSs mean that the server rack has two or more uninterruptible power supplies, which are devices that provide backup power and surge protection in case of a power outage or fluctuation. Multiple UPSs can be connected to separate circuits, or to different outlets on the same circuit, to increase the reliability and availability of power for the server rack23.

Question: 286

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

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

Answer: C

Explanation:

A single-threaded application is an application that can only execute one task or process at a time. A single-threaded application can only utilize one CPU core, regardless of how many cores are available or assigned to the virtual machine. Therefore, adding more vCPU cores to the VM will not improve the performance of the application, as it will still be limited by the speed and capacity of one core12. To troubleshoot this issue, the administrator should check if the application is single-threaded or

multi-threaded. This can be done by using tools such as Task Manager, Performance Monitor, or Process Explorer on Windows, or top, htop, or ps on Linux34. If the application is single-threaded, the administrator should consider the following options:

Reduce the number of vCPU cores on the VM to match the number of threads that the application can use. This can help avoid CPU contention and co-stop issues that may arise from having too many vCPUsrelative to the number of physical cores on the host5.

Upgrade the physical CPU on the host to a faster or newer model that can provide higher clock speed and performance for the single core that the application uses.

Optimize the application code or configuration to make it more efficient or multi-threaded, if possible. This can help the application take advantage of multiple cores and improve its performance.

Question: 287

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

HTTP

A. FTP
B. SCP
C. USB

Answer: C

Explanation:

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

Question: 288

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

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

Answer: D

Explanation:

The BIOS (Basic Input/Output System) or UEFI (Unified Extensible Firmware Interface) setup is a program that allows users to configure the hardware settings of a computer, such as the CPU, memory, disk, and boot options. The BIOS/UEFI setup can be accessed by pressing aspecific key (such as F2, F10, or Delete) during the boot process, before the operating system loads12.

One of the settings that can be changed in the BIOS/UEFI setup is the CPU hyperthreading option. Hyperthreading is a technology that enables a single physical CPU core to execute two threads or tasks simultaneously, improving the performance and efficiency of multi-threaded applications. However, some software licenses may limit the number of CPU cores or threads that can be used, and therefore require disabling hyperthreading on the server34.

To disable hyperthreading on a server, the administrator will likely need to enter the BIOS/UEFI setup and navigate to the processor options menu. There, the administrator will find a setting for Intel ® Hyperthreading Technology or Hyperthreading Function, which can be enabled or disabled. The administrator will need to disable this setting and save the changes. This will turn off hyperthreading on the server and reduce the number of logical CPUs to match the number of physical cores5.

Question: 289

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

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

Answer: B

Explanation:

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

Lower cost: Cloud IaaS can reduce the capital and operational expenses of building and maintaining a physical disaster recovery site, asusers only pay for the resources they use on demand12.

Scalability: Cloud IaaS can offer flexible and elastic scalability of resources, as users can easily provision or deprovision resources according to their needs and workload12.

Availability: Cloud IaaS can ensure high availability and reliability of the application, as users can leverage the cloud provider's redundant andgeographically distributed infrastructure12.

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

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

Question: 290

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

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

Answer: B

Explanation:

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

Question: 291

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 dat
a. Which of the following backup methodologies would best fit this scenario?

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

Answer: D

Explanation:

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

Question: 292

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

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

Answer: A

Explanation:

The user has likely fallen victim to a phishing scam, which is a fraudulent attempt to obtain sensitive information, such as passwords, by disguising as a legitimate entity. The link in the email that required the user to change the account password was probably a fake website that mimicked the company's external database, and captured the user's credentials when they entered them. This could compromise the security and integrity of the company's data, as well as the user's identity and privacy12.

The company should take immediate action to prevent further damage and investigate the incident. The first step is to disable the user's account and inform the security team. Disabling the user's account can prevent unauthorized access to the external database by the attackers, who may use the stolen credentials to log in and manipulate or steal data. Informing the security team can alert them

of the breach and allow them to take appropriate measures, such as scanning for malware, changingpasswords, notifying other users, and reporting the incident34.

Question: 293

Which of the following backup methods protects all the changes that have occurred since the last full backup?

A. Incremental
B. Archive
C. Differential
D. Snapshot

Answer: C

Explanation:

A differential backup is a backup method that protects all the changes that have occurred since the last full backup. A differential backup copies all the files that have been added or modified since the last full backup, regardless of whether they have been backed up before or not. A differential backup does not reset the archive bit of the files, which

means they will be backed up again in the next differential backup. A differentialbackup requires more storage space and time than an incremental backup, but it simplifies the restoration process, as only the last full backup and the last differential backup are needed to restore the data

Question: 294

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

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

Answer: B

Explanation:

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

Ask the user for the date
echo"Enter the date (YYYY-MM-DD):"
readdate

Loop through all the files in the current directory forfilein*
do
Check if the file was created before the date
if[[$(date-r"$file"+%F) <$date]]
then
Move the file to another location mv"$file"/path/to/destination fi
done
Copy

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

A comparator is a script construct that allows the script to compare two values and determine their relationship. A comparator can be used to check if a file was created before the date, but it cannot repeat the test for all files1

A conditional is a script construct that allows the script to execute different blocks of code based on certain conditions. A

conditional can be used to decide whether to move a file or not, but it cannot iterate over all files1

1: CompTIA Server+ Certification Exam Objectives

Question: 295

Which of the following often-overlooked parts of the asset life cycle can cause the greatest number of issues in relation to Pll exposure?

A. Usage
B. End-of-life
C. Procurement
D. Disposal

Answer: D

Explanation:

Disposal is the part of the asset life cycle that can cause the greatest number of issues in relation to PII exposure. PII stands for personally identifiable information, which is any data that can be used to identify a specific individual, such as name, address, phone number, email, social security number, etc. PII exposure is the unauthorized access or disclosure of PII, which can result in identity theft, fraud, or other harms to the individuals whose data is compromised. Disposal is the process of getting rid of an asset that is no longer needed or useful, such as a server, a hard drive, or a mobile device. If the disposal is not done properly, the PII stored on the asset may still be accessible or recoverable by unauthorized parties, such as hackers, thieves, or competitors. Therefore, it is important to follow best practices for secure disposal of assets that contain PII, such as wiping, encrypting, shredding, or physically destroying the data storage media

Question: 296

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

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

Answer: C

Explanation:

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

Question: 297

An administrator needs to reconfigure a teamed network connection on a server in a remote data center. Which of the following will offer the most resilient connection while performing this change?

A. Use of an 00B solution B. Use of a crash cart C. Use of a VNC console D. Use of an RDP console

Answer: A

Explanation:

An out-of-band (OOB) solution is a method of accessing and managing a server remotely without using the network connection or the operating system of the server. An OOB solution can use a dedicated management port, a serial console, or a KVM switch to provide a resilient connection while performing changes to the network configuration of the server. An OOB solution is more reliable than a VNC or RDP console, which depend on the network and the operating system, and more convenient than a crash cart, which requires physical access to the server.
Reference: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 2, Lesson 2.3, Objective 2.3

Question: 298

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

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

Answer: D

Explanation:

A kernel upgrade is a process of updating the core component of a Linux operating system that manages the hardware, memory, processes, and drivers. A kernel upgrade can improve the performance, security, and compatibility of the system, but it can also introduce errors if some modules are not compatible with the new kernel version. Modules are pieces of code that can be loaded and unloaded into the kernel to provide additional functionality or support for specific devices. If a module is not compatible with the kernel, it can cause crashes or instability. Reference: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 4, Lesson 4.2, Objective 4.2

Question: 299

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

Question: 300

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 andmanage 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

Question: 301

A server administrator has been asked to implement a password policy that will help mitigate the chance of a successful brute-force attack. Which of

the following password policies should the administrator implement first?

A. Lockout
B. Length
C. Complexity
D. Minimum age

Answer: B

Explanation:

Password length is the first password policy that the administrator should implement to help mitigate the chance of a successful brute-force attack. A brute-force attack is a method of guessing passwords by trying all possible combinations of characters untilthe correct one is found. The longer the password, the more combinations there are, and the more time and resources it takes to crack it. Therefore, password length is a key factor in password strength and security.
Reference: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 3, Lesson 3.2, Objective 3.2

Question: 302

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

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

Answer: A

Explanation:

The most likely cause of the issue is that the disk uses GPT.GPT stands for GUID Partition Table, which is a newer standard for disk partitioning that supports larger disks and more partitions than theolder MBR (Master Boot Record) standard1.However, GPT is not compatible with some older operating systems, such as Windows XP or Windows Server 20032. Therefore, if the data drive was formatted with GPT on a new Windows server and then moved to an older Windows server, the older server may not be able to recognize the GPT partitions and access the data on the drive. The partition being formatted with ext4, FAT32, or MBR are not likely causes of the issue.Ext4 is a file system that is commonly used on Linux-based systems, but it can also be read by Windows with some third-party software3.FAT32 is a file system that is widely compatible with most operating systems and devices, but it has some limitations such as a maximum file size of 4 GB and a maximum partition size of 8 TB4.MBR is not a file system, but a partitioning scheme that can support various file systems such as NTFS, FAT32, or exFAT5.However, MBR has some disadvantages compared to GPT, such as a maximum disk size of 2 TB and a maximum number of primary partitions of four1.

Question: 303

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

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

Answer: C

Explanation:

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

Question: 304

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: B,F

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.

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

Question: 305

An organization stores backup tapes of its servers at cold sites. The organization wants to ensure the tapes are properly maintained and usable during a DR scenario. Which of the following actions should the organization perform?

A. Have the facility inspect and inventory the tapes on a regular basis.
B. Have duplicate equipment available at the cold site.
C. Retrieve the tapes from the cold site and test them.
D. Use the test equipment at the cold site to read the tapes.

Answer: C

Explanation:

The organization should retrieve the tapes from the cold site and test them to ensure they are properly maintained and usable during a DR scenario. A cold site is a location that has space and power for backup equipment, but no actual equipment installed or configured. The organization stores backup tapes of its servers at cold sites as a precaution in case of a disaster that affects its primary site. However, backup tapes can degrade over time due to environmental factors such as temperature, humidity, dust, or magnetic fields. Therefore, the organization should periodically retrieve the tapes from the cold site and test them on compatible equipment to verify their integrity and readability.

Reference: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 6, Lesson 6.4, Objective 6.4

Question: 306

A server administrator needs to ensure all Window-based servers within a data center have RDP disabled. There are

thousands of servers performing various roles. Which of the following is the best way to meet this requirement?

A. Run chkconfig ——1eve1 345 RDP off.
B. Create a PowerShell script to disable the RDP service.
C. Run chkconfig —— list RDP.
D. Create a Bash shell script to disable the Windows Remote Management service.
E. Create a GPO to disable the Windows Remote Management service.

Answer: B

Explanation:

The best way to meet this requirement is to create a PowerShell script to disable the RDP service on all Windows-based servers within a data center. PowerShell is a scripting language and commandline tool that can be used to automate tasks and manage Windows systems remotely. A PowerShell script can use cmdlets (commands) and parameters to perform actions on multiple servers at once, such as disabling a service or changing a configuration setting. RDP (Remote Desktop Protocol) is a service that allows remote access and control of a Windows system through a graphical user interface. Disabling RDP can improve security by preventing unauthorized or malicious access to the servers.

Reference: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 4, Lesson 4.3, Objective 4.3; Chapter 7, Lesson 7.1, Objective 7.1

Question: 307

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 itprovides exactly two usable IP addresses withoutwasting any addresses or creating any conflicts1.

Question: 308

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

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

Answer: D

Explanation:

The solution the administrator implemented extended runtime. Runtime is the amount of time that a UPS can provide backup power to a server in case of a power outage. By installing additional batteries in the UPS, the administrator increased the capacity and duration of the backup power, allowing the server more time to shut down properly.

Reference: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 1, Lesson 1.4, Objective 1.4

Question: 309

Which of the following concepts is in use when dual power supplies are connected to different power sources?

A. Fault tolerance
B. Active-passive
C. Component redundancy
D. Heartbeat
E. Link aggregation

Answer: A

Explanation:

The concept in use when dual power supplies are connected to different power sources is fault tolerance. Fault tolerance is the ability of a system to continue operating without interruption or loss of data in the event of a failure of one or more components. By connecting dual power supplies to different power sources, the system can switch to the alternative power supply or source if one fails, ensuring continuous availability and reliability.

Reference: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 1, Lesson 1.3, Objective 1.3

Question: 310

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

A. 10.3.7.27
B. 127.0.0.1
C. 192.168.7.1
D. 216,176,128.10

Answer: D

Explanation:

The IP address that best fits this scenario is 216.176.128.10. This is a public IP address that belongs to a range of addresses that are assigned and registered by an Internet service provider (ISP) and can be accessed from anywhere on the Internet. The company has a public range of IP addresses and uses them internally, which means that they do not use private IP addresses or network address translation (NAT) to communicate within their network.
Reference: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 2, Lesson 2.2, Objective 2.2

Question: 311

Which of the following is the most effective way to mitigate risks associated with privacy-related data leaks when sharing with a third party?

A. Third-party acceptable use policy
B. Customer data encryption and masking
C. Non-disclosure and indemnity agreements
D. Service- and operational-level agreements

Answer: B

Explanation:

The most effective way to mitigate risks associated with privacy-related data leaks when sharing with a third party is customer data encryption and masking. Encryption is a process of transforming data into an unreadable format that can only be decrypted with a key or password. Masking is a process of hiding or replacing sensitive data with fake or meaningless data. By encrypting and masking customer data, the organization can protect the confidentiality and integrity of the data and prevent unauthorized access or disclosure by the third party.
Reference: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 3, Lesson 3.3, Objective 3.3

Question: 312

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

Question: 313

Which of the following is used for fail over, providing access to all the services currently in use by an organization without having to physically move any servers or employees?

A. The cloud
B. A cold site
C. A warm site
D. An emergency operations center

Answer: A

Explanation:

The solution that is used for failover, providing access to all the services currently in use by an organization without having to physically move any servers or employees, is the cloud. The cloud is a term that refers to a network of remote servers that are hosted on the Internet and provide various services, such as storage, computing, networking, and applications. The cloud can be used for failover, which is a backup operation that automatically switches to a standby system or service in case of a failure or disruption of the primary system or service. By using the cloud for failover, an organization can ensure continuous availability and accessibility of its services without requiring any physical relocation or intervention.

Reference: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 6, Lesson 6.4, Objective 6.4

Question: 314

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

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

Answer: A

Explanation:

The answer to this question may depend on several factors, such as the number and type of Windows VMs, the number and type of host machines, the number and type of users, and the specific licensing terms and conditions of each licensing model. However, based on the information available from the web search results, one possible answer is per user.

Per user licensing model is a licensing model that allows a user to access Windows VMs from any device, regardless of the number of devices or VMs. Per user licensing model is available for Windows 10 Enterprise E3/E5, Windows VDA E3/E5, and Microsoft 365 F3/E3/E5. Per user licensing model may offer the greatest number of concurrent Windows VMs to run on a host for the lowest cost if the following conditions are met:

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

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

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

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

According to the web search results1, per user licensing model costs $84 per user per year for Windows 10 Enterprise E3, $168 per user per year for Windows 10 Enterprise E5, $100.80 per user per year for Windows VDA E3, and $196.80 per user per year for Windows VDA E5.These prices are based on the Open License Program and may vary depending on the volume and agreement level2 Per core licensing model is a licensing model that requires a license for each core of the processor on the host machine that runs Windows VMs. Per core licensing model is available for Windows Server 2022 Datacenter andStandard editions. Per core licensing model may offer a lower cost than per user licensing model if the following conditions are met:

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

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

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

According to the web search results2, per core licensing model costs $6,155 for 16 core licenses for Windows Server 2022 Datacenter edition and $1,069 for 16 core licenses for Windows Server 2022 Standard edition.These pricesare suggested retail prices and may vary depending on the reseller2 Per instance licensing model is a licensing model that requires a license for each instance of Windows that runs on a host machine or a VM. Per instance licensing model is available for Windows Server 2022 Essentials edition and some older versions of Windows Server. Per instance licensing model may offer a lower cost than per user or per core licensing model if the following conditions are met: The host machine runs only one instance of Windows Server with low resource consumption.

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

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

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

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

Question: 315

Users are able to connect to the wireless network, but they are unable to access the internet. The network administrator verifies connectivity to all network devices, and there are no ISP outages. The server administrator removes the old address leases from the active leases pool, which allows users to access the internet. Which of the following is most likely causing the internet issue?

A. THe DHCP exclusion needs to be removed.
B. The DHCP scope is full.
C. The DHCP scope options are misconfigured.
D. The DHCP lease times are too short.
E. The DHCP reservations need to be configured.

Answer: B

Explanation:

The most likely cause of the internet issue is B. The DHCP scope is full.

A DHCP scope is a range of IP addresses that a DHCP server can assign to DHCP clients on a network. A DHCP scope has a start address and an end address, and it can also have some excluded addresses that are not available for lease. A DHCP scope can have various options, such as subnet mask, default gateway, DNS server, etc., that are applied to the DHCP clients along with the IP address. A DHCP scope also has a lease time, which is the duration that a DHCP client can use an IP address before renewing it or releasing it. A DHCP scope can have reservations, which are fixed IP addresses that are assignedto specific DHCP clients based on their MAC addresses12

If a DHCP scope is full, it means that there are no more IP addresses available for lease in the scope. This can happen if the number of DHCP clients exceeds the number of IP addresses in the scope, or if the lease time is too long and the IP addresses are not released or reused frequently enough. If a DHCP scope is full, any new or existing DHCP clients that request an IP address from the DHCP server will not receive one, and they will not be able to access the network or the internet12

In this scenario, users are able to connect to the wireless network, but they are unable to access the internet. The network administrator verifies connectivity to all network devices, and there are no ISP outages. The server administrator removes the old address leases from the active leases pool, which allows users to access the internet. This indicates that the DHCP scope is full, and that removing the old leases frees up some IP addresses for lease in the scope. Therefore, option B is the most likely cause of the internet issue.

Question: 316

A server administrator is building a pair of new storage servers. The servers will replicate; therefore, no redundancy is required, but usable capacity must be maximized. Which of the following RAID levels should the server administrator implement?

A. 0
B. 1
C. 5
D. 6
E. 10

Answer: A

Explanation:

The RAID level that should be implemented to maximize usable capacity without requiring redundancy is RAID 0. RAID (Redundant Array of Independent Disks) is a technology that combines multiple physical disks into a logical unit that

provides improved performance, reliability, or both. RAID 0 is a RAID level that splits data evenly across two or more disks without parity or mirroring. RAID 0 does not provide any redundancy or fault tolerance, but it increases usable capacity and performance by allowing parallel read and write operations.
Reference: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 1, Lesson 1.2, Objective 1.2

Question: 317

An administrator is working on improving the security of a new domain controller. A report indicates several open ports on the server. Which of the following ports should the administrator disable?

A. 135
B. 636
C. 3268
D. 3389

Answer: D

Explanation:

The port that should be disabled on the firewall is port 3389. Port 3389 is used by Remote Desktop Protocol (RDP), which is a protocol that allows remote access and control of a Windows system through a graphical user interface. RDP can pose a security risk if it is not properly configured or secured, as it can expose the system to unauthorized or malicious access from external sources. Therefore, port 3389 should be disabled on the firewall unless it is needed for legitimate purposes.
Reference: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 3, Lesson 3.3, Objective 3.3

Question: 318

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

A. 25
B. 443
C. 3389

D. 8080

Answer: B

Explanation:

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

Question: 319

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

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

Answer: A

Explanation:

A method of application hardening is installing the latest patches. Application hardening is a process of reducing the attack surface and vulnerabilities of an application by applying security measures and best practices. Installing the latest patches is one way to harden an application, as patches are updates that fix bugs, errors, or security issues in an application. By installing the latest patches, an application can be protected from known exploits or threats.
Reference: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 3, Lesson 3.5, Objective 3.5

Topic 3, Exam Set C

Question: 320

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

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

Answer: C

Explanation:

The first thing that the administrator should check is the network cable. The network cable is a physical medium that connects a server to a switch or other network device. The network cable can affect the speed and quality of the network connection, depending on its type, length, and condition. If the network cable is damaged, faulty, or incompatible, it can cause the server to connect at a lower speed than expected. Therefore, the administrator should check the network cable for any signs of wear, tear, or mismatch, and replace it if necessary.
Reference: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 2, Lesson 2.1, Objective 2.1

Question: 321

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

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

Answer: B

Explanation:

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

Question: 322

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

Question: 323

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

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

Answer: A

Explanation:

The most likely cause of the issue is that the server has not been joined to the domain. A domain is a logical group of computers and devices that share a common directory service and security policy. A domain controller is a server that manages the domain and authenticates users and computers that want to access domain resources. To log on with domain credentials, a server must be joined to the domain and registered in the directory service. If a server has not been joined to the domain, it will not be recognized or authorized by the domain controller.
Reference: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 4, Lesson 4.3, Objective 4.3

Question: 324

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

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

Answer: B

Explanation:

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

Question: 325

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 bySCCM. 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. Reference: CompTIA Server+ Certification Exam Objectives, Domain 5.0: Security, Objective 5.4: Given a scenario, apply patches/updates and validate their installation.

Question: 326

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

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

Answer: D

Explanation:

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

Given a scenario, install and deploy primary storage devices based on given specifications and interfaces.

Question: 327

IDS alerts indicate abnormal traffic patterns are coming from a specific server in a data center that hosts sensitive data. Upon further investigation, the server administrator notices this server has been infected with a virus due to an exploit of a known vulnerability from its database software. Which of the following should the administrator perform after removing the virus to mitigate this issue from reoccurring and to maintain high availability? (Select three).

A. Run a vulnerability scanner on the server.
B. Repartition the hard drive that houses the database.
C. Patch the vulnerability.
D. Enable a host firewall.
E. Reformat the OS on the server.
F. Update the antivirus software.
G. Remove the database software.
H. Air gap the server from the network.

Answer: A,C,F

Explanation:

After removing the virus from the server, the administrator should perform the following actions to mitigate the issue from reoccurring and to maintain high availability:

Run a vulnerability scanner on the server to identify any other potential weaknesses or exposures that could be exploited by attackers.

Patch the vulnerability that allowed the virus to infect the server in the first place, using the latest updates from the database software vendor or a trusted source.

Update the antivirus software on the server to ensure it has the most recent virus definitions and can detect and prevent future infections. The other options are either unnecessary or counterproductive for this scenario. Repartitioning the hard drive, reformatting the OS, removing the database software, or air gapping the server from the network would cause downtime and data loss, while enabling a host firewall would not prevent a virus infection from within the network. Reference: CompTIA Server+ Certification Exam Objectives, Domain 5.0: Security, Objective 5.2: Given a scenario involving a security threat/vulnerability/risk, implement appropriate mitigation techniques.

Question: 328

A server technician notices several of the servers in a data center are making loud noises. The servers are still working correctly, and no indicator lights show any issues. Which of the following should the technician do first to ensure the issues are corrected and the servers remain online?

A. Replace the drives.
B. Upgrade the firmware.
C. Establish a remote connection to the server.
D. Replace the fans.

Answer: A

Explanation:

The loud noises from the servers are most likely caused by failing hard disk drives, which can produce clicking or grinding sounds. Replacing the drives with new ones can prevent data loss and downtime. Replacing the drives can be done without shutting down the server if they are hot-swappable, which means they can be removed and inserted while the server is running. Reference: CompTIA Server+ Certification Exam Objectives, Domain 3.0: Storage, Objective 3.1: Given a scenario, install, deploy, configure and update physical storage devices.

Question: 329

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

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

Answer: A

Explanation:

An access control vestibule is a physical security concept that limits personnel access to a restricted area within a data center. It is a small room or hallway that has two doors: one that leads to the outside and one that leads to the restricted area. The doors are controlled by an electronic lock that requires authentication, such as a card reader, biometric scanner, or keypad. Only authorized

personnel can enter the vestibule and access the restricted area. Reference: CompTIA Server+ Certification Exam Objectives, Domain 5.0: Security, Objective 5.1: Given a scenario, apply physical security methods to a server.

Question: 330

A technician is sizing a new server and, for service reasons, needs as many hot-swappable components as possible. Which of the following server components can most commonly be replaced without downtime? (Select three).

A. Drives

B. Fans
C. CMOSIC
D. Processor
E. Power supplies
F. Motherboard
G. Memory
H. BIOS

Answer: A,B,E

Explanation:

Drives, fans, and power supplies are server components that can most commonly be replaced without downtime if they are hot-swappable. Hot-swappable components can be removed and inserted while the server is running, without affecting its operation or performance. Drives store data and applications, fans cool down the server components, and power supplies provide electricity to the server. Replacing these components can prevent data loss, overheating, or power failure.
Reference: CompTIA Server+ Certification Exam Objectives, Domain 2.0: Hardware, Objective 2.2: Given a scenario, install, configure and maintain server components.

Question: 331

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

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

Answer: D

Explanation:

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

Question: 332

A new 40GB NIC has just been installed in a server but is not detected within the Windows server OS.
Which of the following would most likely fix the issue?

A. Update the firmware on the NIC.
B. Update the server OS.
C. Update the remote management console.
D. Update the switch firmware.

Answer: A

Explanation:

Updating the firmware on the NIC is the most likely solution to fix the issue of a new 40GB NIC not being detected within the Windows server OS. Firmware is a softwareprogram that controls the functionality of a hardware device, such as a NIC (network interface card). A NIC is a device that enables network communication for a server by providing an interface between the server and the network cable or wireless connection. Updating the firmware on the NIC can improve its performance, compatibility, and stability with the server OS and network protocols. Reference: CompTIAServer+ Certification Exam Objectives, Domain 4.0: Networking, Objective 4.1: Given a scenario, configure network settings for servers.

Question: 333

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

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

Answer: D

Explanation:

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

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

Reference: CompTIA Server+ Certification Exam Objectives, Domain 6.0: Troubleshooting, Objective 6.3: Given a scenario,

troubleshoot scripting errors using PowerShell commands.

Question: 334

A Linux server requires repetitive tasks for reconfiguration. Which of the following would be the best scripting language to use?

A. PowerShell
B. Batch command file
C. Bash
D. Visual Basic

Answer: C

Explanation:

Bash is a scripting language that is commonly used in Linux systems to automate tasks and manipulate text. Bash scripts can run commands, variables, functions, loops, and conditional statements. PowerShell is a scripting language that is mainly used in Windows systems, while batch command files are simple text files that contain a series of commands to be executed by the command-line interpreter. Visual Basic is a programming language that is used to create applications, not scripts. Reference: [CompTIA Server+ Certification Exam Objectives], Domain 4.0: Server Administration, Objective 4.2: Given a scenario, perform proper server maintenance techniques.

Question: 335

An administrator reviews a new server that was received from a vendor and notes the OS has been installed to a two-drive array configured with RAID 0. Which of the following best describes what will happen if a drive in that array fails?

A. The server will gracefully shut down.
B. The server will immediately crash.
C. The server will operate but in read-only mode.
D. The server will continue to operate normally.

Answer: B

Explanation:

RAID 0 is a configuration that splits data evenly across two or more disks without parity or mirroring. This improves performance but offers no fault tolerance. If a drive in a RAID 0 array fails, the data on the array becomes inaccessible and the server will immediately crash. The other options are not applicable to RAID 0. Reference: [CompTIA Server+ Certification Exam Objectives], Domain 3.0: Storage, Objective 3.2: Given a scenario, configure RAID using best practices.

Question: 336

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

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

Answer: A

Explanation:

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

Question: 337

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

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

Answer: D

Explanation:

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

Question: 338

Which of the following security risks provides unauthorized access to an application?

A. Backdoor

B. Data corruption

C. Insider threat

D. Social engineering

Answer: A

Explanation:

A backdoor is a security risk that provides unauthorized access to an application. A backdoor is a hidden or undocumented way of bypassing the normal authentication or encryption mechanisms of an application, allowing an attacker to gain remote access, execute commands, or steal data. A backdoor can be created intentionally by the developer, maliciously by an attacker, or unintentionally by a programming error. Reference: CompTIA Server+ Certification Exam Objectives, Domain 5.0: Security, Objective 5.2: Given a scenario, apply logical access control methods.

Question: 339

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

A. Close ports 69 and 1010 and rerun the scan.

B. Close ports 80 and 443 and rerun the scan.

C. Close port 3389 and rerun the scan.

D. Close all ports and rerun the scan.

Answer: C

Explanation:

The server administrator should close port 3389 and rerun the scan. Port 3389 is used for Remote Desktop Protocol (RDP), which allows remote access and control of a server. RDP is vulnerable to brute-force attacks, credential theft, and malware infection. Closing port 3389 can prevent unauthorized access and improve the security of the server. The other ports are not as risky as port 3389 and can be left open for legitimate purposes. Reference: CompTIA Server+ Certification Exam Objectives, Domain 5.0: Security, Objective 5.4: Given a scenario, implement proper environmental controls and techniques.

Question: 340

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

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

Answer: C

Explanation:

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

Question: 341

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

A. Definition updates

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

Answer: B

Explanation:

Driver updates should be applied to the new components that have been added to a mission-critical server, as part of the server hardening requirements. Drivers are software programs that enable the communication and functionality of hardware devices, such as network cards, storage controllers, or

graphics cards. Updating drivers can improve the performance, compatibility, and stability of the new components with the server operating system and applications. Reference: CompTIA Server+ Certification Exam Objectives, Domain 2.0: Hardware, Objective 2.2: Given a scenario, install, configure and maintain server components.

Question: 342

A company created a new DR plan. The management team would like to begin performing a review of this plan without endangering company data and with a minimal time commitment. Which of the following testing methods would best allow for this type of review?

A. Simulated
B. Tabletop
C. Live
D. Non-production

Answer: B

Explanation:

Tabletop testing is a method of reviewing a DR plan without endangering company data and with a minimal time commitment. Tabletop testing involves a simulated scenario where the participants discuss their roles and responsibilities, identify potential issues, and evaluate the effectiveness of the plan. Simulated, live, and non-production testing are methods that involve more time and resources, and may pose some risks to company data. Reference: [CompTIA Server+ Certification Exam Objectives], Domain 6.0: Disaster Recovery, Objective 6.3: Compare and contrast various backup techniques.

Question: 343

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

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

Answer: B

Explanation:

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

Question: 344

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

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

Answer: A

Explanation:

Drive encryption is a form of encryption that best achieves the goal of data-at-rest encryption for all data. Drive encryption encrypts the entire hard drive, including the operating system, applications, and files. This prevents unauthorized access to the data if the drive is lost or stolen. Database, folder, and file encryption are forms of encryption that only encrypt specific data sets, not all
data. Reference: [CompTIA Server+ Certification Exam Objectives], Domain 5.0: Security, Objective 5.3: Given a scenario involving a security threat/vulnerability/risk, implement appropriate mitigation techniques.

Question: 345

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

A. Warranty
B. Purchase order
C. License
D. Baseline document

Answer: A

Explanation:

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

Question: 346

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

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

Answer: B

Explanation:

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

involving network security/access methods, implement an appropriate solution.

Question: 347

Which of the following attacks is the most difficult to mitigate with technology?

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

Answer: D

Explanation:

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

Question: 348

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

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

Answer: C

Explanation:

A public cloud model will most likely meet the need of eliminating all company-owned data centers.

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

Question: 349

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

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

Answer: A

Explanation:

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

Question: 350

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

A. Inform the impacted users.
B. Make the changes to the system.
C. Determine the probable causes.

D. Identify changes to the server.

Answer: C

Explanation:

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

Identify the problem: Gather information from various sources, such as users, logs, or alerts, to understand the symptoms and scope of the problem.

Establish a theory of probable cause: Analyze the information and formulate one or more possible causes of the problem based on evidence or experience.

Test the theory to determine cause: Perform tests or experiments to verify or eliminate each possible cause until the root cause is found.

Establish a plan of action to resolve the problem and implement the solution: Design and execute a plan to fix the problem using appropriate tools and techniques.

Verify full system functionality and implement preventive measures: Confirm that the problem is resolved and that no other issues arise as a result of the solution. Implement preventive measures to avoid recurrence of the problem or improve performance.

Document findings, actions, and outcomes: Record the details of the problem, its cause, its solution, and its outcome for future reference or knowledge sharing. Reference: [CompTIA Server+ Certification Exam Objectives], Domain 6.0: Troubleshooting, Objective 6.1: Given a scenario involving server hardware issues (e.g., power supply failure), troubleshoot using appropriate tools.

Question: 351

An upper management team is investigating a security breach of the company's filesystem. It has been determined that the breach occurred within the human resources department. Which of the following was used to identify the breach in the human resources department?

A. User groups
B. User activity reports
C. Password policy
D. Multifactor authentication

Answer: B

Explanation:

User activity reports were used to identify the security breach in the human resources department. User activity reports are

records of the actions and events performed by users on a system or network, such as login/logout times, files accessed or modified, commands executed, or websites visited. User activity reports can help monitor and audit user behavior, detect and investigate security incidents, and enforce policies and compliance. User activity reports can be generated by various tools, such as log management software, security information and event management (SIEM) systems, or user and entity behavior analytics (UEBA) solutions. Reference: [CompTIAServer+ Certification Exam Objectives], Domain 5.0: Security, Objective 5.2: Given a scenario, apply logical access control methods.

Question: 352

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

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

Answer: A

Explanation:

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

Question: 353

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

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

Answer: A

Explanation:

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

Question: 354

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

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

Answer: C

Explanation:

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

Question: 355

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

A. RAIDO
B. RAID1
C. RAID 5
D. RAID 6
E. RAID 10

Answer: E

Explanation:

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

Question: 356

A server technician is placing a newly configured server into a corporate environment. The server will

be used by members of the accounting department, who are currently assigned by the VLAN identified below:

VLAN name	VLAN ID	IP address	Default gateway	Exclusion range
Accounting	25	172.16.25.1 172.16.25.254/24	172.16.25.254	172.16.25.50 172.16.25.100

Which of the following IP address configurations should the technician assign to the new server so the members of the accounting group can access the server?

A. IP address: 172.16.25.90/24 Default gateway: 172.16.25.254
B. IP address: 172.16.25.101/16 Default gateway: 172.16.25.254
C. IP address: 172.16.25.254/24 Default gateway: 172.16.25.1
D. IP address: 172.16.26.101/24 Default gateway: 172.16.25.254

Answer: A

Explanation:

The IP address configuration that the technician should assign to the new server so the members of the accounting group can access the server is 172.16.25.90/24 for the IP address and 172.16.25.254 for the default gateway. This configuration matches the VLAN identified in the image, which has a network address of 172.16.25.0/24 and a subnet mask of 255.255.255.0. The IP address of the server must be within the same network range as the VLAN, which is from 172.16.25.1 to 172.16.25.254, excluding the network and broadcast addresses (172.16.25.0 and 172.16.25.255). The default gateway of the server must be the same as the VLAN, which is 172.16.25.254, to allow communication with other networks or devices outside the VLAN. Reference: [CompTIA Server+ Certification Exam Objectives], Domain 4.0: Networking, Objective 4.1: Given a scenario, configure network settings for servers.

Question: 357

An application server cannot communicate with a newly installed database server. The database server, which has static IP information, is reading the following output from ipconf ig:

The application server is reading the following output from ipconf ig:

Which of the following most likely contains an error?

A. IP address
B. DHCP
C. Gateway
D. Subnet mask

Answer: D

Explanation:

The subnet mask is most likely containing an error that prevents the application server from communicating with the newly installed database server. The 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 subnet mask determines which devices belong to the same network or subnet and can communicate directly with each other without routing or switching devices. The subnet mask of the database server is 255.255.O.O, which means that all 32 bits of its IP address are used for the network portion and none for the host portion, which is invalid and makes it unreachable by any other device on any network or subnet. The subnet mask of the application server is 255.O.O.O, which means that only 8 bits of its IP address are used for the network portion and 24 bits are used for the host portion, which is also uncommon and makes it incompatible with most networks or subnets. Reference: [CompTIA Server+ Certification Exam Objectives], Domain 4.0: Networking, Objective 4.1: Given a scenario, configure network settings for servers.

Question: 358

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. Reference: [CompTIA Server+ Certification Exam Objectives], Domain 5.0: Security, Objective 5.3: Given a scenario, explain methods and techniques to secure data.

Question: 359

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

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

Answer: D

Explanation:

The most likely reason for this result is that there is a manual entry in the hosts file that maps the FQDN

to an incorrect IP address (104.18.17.32). The hosts file is a text file that contains mappings of hostnames or domain names to IP addresses, which are used by the operating system to resolve names before querying DNS servers on the network or internet. The hosts file can be used to override DNS settings or block access to certain websites by redirecting them to different IP addresses, such as localhost (127.O.O.l) or invalid addresses (O.O.O.O). If there is a manual entry in the hosts file that conflicts with DNS records, it can cause name resolution errors or connectivity issues. Reference: [CompTIA Server+ Certification Exam Objectives], Domain 4.0: Networking, Objective 4.l: Given a scenario, configure network settings for servers.

Question: 360

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

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

Answer: D

Explanation:

The most likely cause of this issue is that the printer had DHCP enabled, which changed its IP address after adding it to the print server and configuring it on each user's machine. DHCP (Dynamic Host Configuration Protocol) is a network protocol that assigns IP addresses and other network configuration parameters to devices automatically, without manual intervention. DHCP can simplify network management and avoid IP conflicts, but it can also cause problems if the devices are not configured to use static or reserved IP addresses. If the printer had DHCP enabled, it might have received a different IP address from the DHCP server after rebooting or reconnecting to the network, which would make it unreachable by the print server and the users' machines that were configured with the previous IP address. Scanning to email would still work, as it does not depend on the print server or the users' machines, but on the printer's SMTP settings and internet connection. Reference: [CompTIA Server+ Certification Exam Objectives], Domain 4.0: Networking, Objective 4.l: Given a scenario, configure network settings for servers.

Question: 361

A server administrator is setting up a disk with enforcement policies on how much data each home share can hold. The amount of data that is redundant on the server must also be minimized. Which of the following should the administrator perform on the server? (Select two).

A. Partitioning
B. Deduplication
C. Disk quotas
D. Compression
E. Cloning
F. Provisioning

Answer: B,C

Explanation:

Deduplication is a process that eliminates redundant data blocks and reduces the amount of storage space needed. Disk quotas are policies that limit the amount of disk space that each user or group can use on a volume.

Reference:

CompTIA Server+ Certification Exam Objectives1, page 8

Data Deduplication interoperability2

Question: 362

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

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

Answer: B,E

Explanation:

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

Reference:

CompTIA Server+ Certification Exam Objectives1, page 9

Signs or symptoms of a CMOS battery failure2

NTP: Network Time Protocol

Question: 363

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

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

Answer: C

Explanation:

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

Reference:

CompTIA Server+ Certification Exam Objectives1, page 8

RAID Levels and Types Explained: Advantages and Disadvantages2

RAID Levels & Fault Tolerance3

Question: 364

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

A. 53
B. 80
C. 389
D. 443
E. 445

F. 3389
G. 8080

Answer: D,F

Explanation:

Port 443 is used for HTTPS, which is a secure version of HTTP that encrypts the data between the web server and the client. Port 3389 is used for RDP, which is a protocol that allows remote desktop connections to a server. These ports should be allowed for a secure web server that can be administered via RDP.
Reference:
CompTIAServer+ Certification Exam Objectives1, page 15
Common Ports Cheat Sheet: The Ultimate Ports & Protocols List2

Question: 365

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.

Reference:

CompTIA Server+ Certification Exam Objectives1, page 14

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

Question: 366

A server administrator is reviewing the following specifications:

VM01 Host:

CPU: 2 Physical, 4 Cores

RAM: 16GB

Storage: 16TB

Server 1 on VM01:

CPU: 1 virtual socket, 1 core per socket

RAM: 4GB

Storage: 8TB

Server 2 on VM01:

CPU: 2 virtual sockets, 2 cores per socket

RAM: 8GB

Storage: 10TB

Which of the following is described given these specifications?

A. Virtual switch
B. Host vs. guest
C. Overprovisioning
D. Scalability

Answer: C

Explanation:

Overprovisioning is a situation where the allocated resources for a virtual machine exceed the available resources of the physical host. In this case, the storage allocated for Server 1 and Server 2 on VM01 is 8TB and 10TB respectively, which adds up to 18TB. However, the storage available on the VM01 host is only 16TB, which means there is a 2TB deficit. This can cause performance issues and errors for the virtual machines.

Reference:

CompTIA Server+ Certification Exam Objectives1, page 8

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

Question: 367

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

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

Answer: B,E

Explanation:

Archive is a method of storing historical data that is not frequently accessed or modified. Archive can reduce the amount of data that needs to be backed up and save bandwidth and storage space. Differential is a method of backing up only the data that has changed since the last full backup. Differential can also save bandwidth and storage space, as well as speed up the backup process. Reference:
CompTIA Server+ Certification Exam Objectives1, page 12
Server Management: Server Hardware Installation and Management2, Module 2, Lesson 5

Question: 368

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

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

Answer: D

Explanation:

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

Reference:

CompTIA Server+ Certification Exam Objectives1, page 15

CompTIA Server+: Authentication & Authorization2

Question: 369

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

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

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

Answer: D

Explanation:

Question: 370

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

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

Answer: A

Explanation:

Application consistent backup is a method of backing up data that ensures the integrity and consistency of the application state. It involves notifying the application to flush its data from memory to disk and quiescing any write operations before taking a snapshot of the data. If the databases were not backed up to be application consistent, they might contain incomplete or corrupted data that cannot be restored properly.

Reference:

CompTIA Server+ Certification Exam Objectives1, page 12

What is Application Consistent Backup and How to Achieve It2

Application-Consistent Backups3

Question: 371

An administrator notices high traffic on a certain subnet and would like to identify the source of the traffic. Which of the following tools should the administrator utilize?

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

Answer: D

Explanation:

Application consistent backup is a method of backing up data that ensures the integrity and consistency of the application state. It involves notifying the application to flush its data from memory to disk and quiescing any write operations before taking a snapshot of the data. If the databases were not backed up to be application consistent, they might contain incomplete or corrupted data that cannot be restored properly.

Reference:

CompTIA Server+ Certification Exam Objectives1, page 12

What is Application Consistent Backup and How to Achieve It2

Application-Consistent Backups3

Question: 372

A site is considered a warm site when it:

has basic technical facilities connected to it.

has faulty air conditioning that is awaiting service.

is almost ready to take over all operations from the primary site.

A. is fully operational and continuously providing services.

Answer: A

Explanation:

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

Question: 373

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

Drive

Database

Folder

A. File

Answer: A

Explanation:

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

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

Question: 374

A company recently implemented VoIP across a multicampus environment with ten locations. The company uses many network technologies, including fiber, copper, and wireless. Users calling between three of the locations have reported that voices sound strange. Which of the following should be monitored to narrow down the issue?

A. Disk IOPS
B. CPU utilization
C. RAM utilization
D. Network latency

Answer: D

Explanation:

Network latency is the measure of delay in data transmission over a network. It can affect the quality of voice over IP (VoIP) calls by causing echo, jitter, or distortion. Network latency can be caused by various factors such as network congestion, distance, routing, or bandwidth. To monitor network latency, you can use tools such as ping, traceroute, or network analyzers.

Reference: CompTIA Server+ Study Guide, Chapter 6: Networking, page 237.

Question: 375

Which of the following distributes a load across all interfaces?

A. Link aggregation group
B. Most recently used algorithm
C. Active-passive configuration
D. Failover

Answer: A

Explanation:

A link aggregation group (LAG) is a technique that combines multiple physical network interfaces into a single logical interface. This allows for the distribution of traffic across all the interfaces in the group, increasing bandwidth and redundancy. A LAG can use different modes to balance the load, such as address hashing, dynamic, or most recently used algorithm.

Reference: CompTIA Server+ Study Guide, Chapter 6: Networking, page 239.

Question: 376

Hackers recently targeted a company with an attack that resulted in a system breach, which compromised the organization's dat

a. 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.

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

Question: 377

After a technician upgrades the firmware on a database server that is connected to two external storage arrays, the server prompts the technician to configure RAID. The technician knows the server had several configured RAID sets and thinks the firmware upgrade cleared the RAID configurations. Which of the following should the technician do to troubleshoot this issue?

A. Power cycle the storage arrays and rescan RAID on the server.
B. Boot the OS into recovery mode and rescan the disks.
C. Restore the default RAID configuration and reboot.
D. Perform a rescan on the server's RAID controller.

Answer: D

Explanation:

A rescan on the server's RAID controller is a possible troubleshooting step to detect the existing RAID configurations on the connected storage arrays. A firmware upgrade may cause the RAID controller to lose the RAID metadata or settings, and a rescan may restore them. A rescan is preferable to restoring the default RAID configuration, as the latter may erase the existing data on the arrays. Power cycling the storage arrays or booting the OS into recovery mode may not help if the RAID controller does not recognize the RAID sets.

Reference: CompTIA Server+ Study Guide, Chapter 7: Storage, page 287.

Question: 378

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

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

Answer: B

Explanation:

A DNS entry is a record that maps a domain name to an IP address. If the DNS entry for the new application is incorrect, users will not be able to resolve the domain name to the correct IP address and port number. This will prevent them from accessing the application, even if the firewall rules allow port 8443. To fix this issue, the server administrator should verify and update the DNS entry for the new application.
Reference: CompTIA Server+ Study Guide, Chapter 6: Networking, page 230.

Question: 379

A startup company needs to set up an initial disaster recovery site. The site must be cost-effective and deployed quickly. Which of the following sites should the company set up?

A. Hot
B. Cold
C. Colocated
D. Warm

Answer: B

Explanation:

A cold site is a backup facility with little or no hardware equipment installed. A cold site is the most cost-effective option among the three disaster recovery sites. However, due to the fact that a cold site doesn't have any pre-installed equipment, it takes a lot of time to properly set it up so as to fully resume business operations1.
Reference = 1: Disaster Recovery Sites Comparison: Which one to Choose? -

NAKIVO(https://www.nakivo.com/blog/overview-disaster-recovery-sites/)

Question: 380

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

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

Answer: B

Explanation:

TFTP (Trivial File Transfer Protocol) is a simple and lightweight protocol that can be used to transfer files over a network. TFTP is often used to upgrade firmware or software on network devices, such as routers, switches, or servers. TFTP can also be used to install a Type 1 hypervisor, such as VMware ESXi, on multiple hosts simultaneously12.

Reference = 1: How to Install VMware ESXi Type 1 Hypervisor - MatthewEaton.net(https://mattheweaton.net/posts/how-to-install-vmware-esxi-type-1-hypervisor/)

2: Explore Type 1 Hypervisors - Set Up Virtual Machines Using VirtualBox and vSphere - OpenClassrooms(https://openclassrooms.com/en/courses/7163136-set-up-virtual-machines-using- virtualbox-and-vsphere/7358546-explore-type-1-hypervisors)

Question: 381

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

A. tracert

B. telnet
C. ipconfig
D. ping

Answer: C

Explanation:

Question: 382

An administrator is troubleshooting connectivity to a remote server. The goal is to remotely connect to the server to make configuration changes. To further troubleshoot, a port scan revealed the ports on the server as follows:

Port 22: Closed
Port 23: Open

Port 990: Closed

Which of the following next steps should the administrator take?

Reboot the workstation and then the server.

A. Open port 990 and close port 23.
B. Open port 22 and close port 23.
C. Open all of the ports listed.
D. Close all of the ports listed.

Answer: B

Explanation:

Port 22 is used for SSH (Secure Shell), which is a secure and encrypted protocol for remote access to a server. Port 23 is used for Telnet, which is an insecure and unencrypted protocol for remote access. Port 990 is used for FTPS (File Transfer Protocol Secure), which is a secure and encrypted protocol for file transfer. The administrator should open port 22 and close port 23 to enable SSH and disable Telnet, as SSH is more secure and reliable than Telnet. The administrator does not need to open port 990, as FTPS is not required for making configuration changes123.

Reference = 1: Remote Desktop - Allow access to your PC from outside your network(https://learn.microsoft.com/en-us/windows-server/remote/remote-desktop- services/clients/remote-desktop-allow-outside-access) 2: Test remote network port connectionin Windows 10 - Winaero(https://winaero.com/test-remote-network-port-connection-in-windows-10/) 3: Windows Command to check if a remote server port is opened?(https://superuser.com/questions/1035018/windows-command-to-check-if-a-remote- server-port-is-opened)

Question: 383

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

Starting Nmap 7.70 https://nmap.org) at 2019-09-19 14:30 UTC
Nmap scan report for www.abc.com (172.45.6.85)
Host is up (0.0021s latency)

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

Port State Service
21/tcp filtered ftp
22/tcp filtered ssh
23/tcp filtered telnet
69/tcp open @username.com
80/tcp open http
110/tcp filtered pop
143/tcp filtered imap
443/tcp open https
1010/tcp open www.popup.com
3389/tcp filtered ms-abc-server

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

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

Answer: A

Explanation:

Port 69 is used for TFTP (Trivial File Transfer Protocol), which is an insecure and unencrypted protocol for file transfer. Port 1010 is used for a malicious website that generates pop-up ads. Both of these ports are likely to be exploited by hackers or malware to compromise the server or the web application. The server administrator should close these ports and rerun the scan to verify that they are no longer open12.

Reference = 1: Why Are Some Network Ports Risky, And How Do You Secure Them? - How-To Geek(https://www.howtogeek.com/devops/why-are-some-ports-risky-and-how-do-you-secure- them/) 2: Switchport Port Security Explained With Examples - ComputerNetworkingNotes(https://www.computernetworkingnotes.com/ccna-study- guide/switchport-port-security-explained-with-examples.html)

Question: 384

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

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

Answer: B

Explanation:

A subscription license is a type of license that grants the user the right to use a product or service for a fixed period of time, usually a year. The user pays a recurring fee, regardless of how much they use the product or service. Subscription licenses are common for cloud-based software and services, such as Microsoft 3651 or DocuSign2.

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

Question: 385

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

this behavior?

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

Answer: A

Explanation:

A server has multiple fans inside the chassis to cool down the components and prevent overheating. If one of the fans fails, it can cause the server to reboot and crash due to thermal issues. A failed fan can also make loud noises due to friction or vibration. The administrator should check the fans and clean them from dust and debris, or replace them if they are damaged12.

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

Question: 386

A technician is installing an OS on ten servers. Which of the following media installation types would allow for the fastest installation time?

A. Network
B. Embedded
C. Optical
D. USB

Answer: A

Explanation:

Network Installation: Allows the OS image to be deployed from a central server, streamlining deployment across multiple systems simultaneously. This is significantly faster than individual installations from other media. (CompTIA Server+ Objectives SK0-004: 3.1)

Why other options are less optimal:

Embedded: Refers to OSes pre-installed on hardware and not intended for mass deployment.
Optical (CDs/DVDs): Requires physical media insertion on each server, slower than network distribution.
USB Similar to optical, requires individual installations and can be time-consuming for multiple servers.

Question: 387

A web server that is being deployed in the perimeter network needs to be shielded from malicious traffic. Which of the following could help identify these threats?

A. Applying OS updates
B. Disabling unused services
C. Implementing HIDS
D. Installing anti-malware

Answer: C

Explanation:

HIDS (Host Intrusion Detection System): Continuously monitors a system for suspicious activity and logs or raises alerts when potential threats are identified. This proactive approach is crucial for identifying and mitigating threats on a web server exposed to the external network.
Applying OS updates: While essential for maintaining system security, updates address vulnerabilities and may not necessarily identify ongoing threats.
Disabling unused services: Reduces the attack surface by minimizing potential entry points for malicious actors, but doesn't actively identify threats.
Installing anti-malware: Primarily designed to detect and remove malware after infection, not for ongoing threat identification.
Reference:
CompTIA Server+ Objectives (Exam codes SK0-004 or SK0-005): Search for sections on intrusion detection and prevention.

Question: 388

A technician is setting up a repurposed server. The minimum requirements are 2TB while ensuring the highest performance and providing support for one drive failure.The technician has the following six drives available:

Which of the following drive selections should the technician utilize to best accomplish this goal?

A. 1,2, 4, and 6
B. 1, 2, 3, 5, and 6
C. 1, 2, 4, 5, and 6
D. 1, 2, 3, 4, and 6

Answer: C

Explanation:

RAID 5 configuration: Using five of the available drives in a RAID 5 configuration meets the requirements for:
Storage capacity: Four 600GB drives (2, 5, and 6) provide a total usable capacity of 2.4TB (4 * 600 * 0.8), exceeding the minimum requirement of 2TB. RAID 5 introduces parity data for fault tolerance, sacrificing some usable space (one drive's worth).
Performance: The combination of multiple drives in a RAID 5 array improves read performance compared to a single drive setup.
Fault tolerance: Even with a single drive failure (any of the five drives used in the RAID 5), the remaining drives can reconstruct the lost data, allowing the server to continue operating.

Question: 389

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

A. 25
B. 53
C. 443
D. 3389
E. 8080

Answer: C

Explanation:

HTTPS (Secure Web Traffic): Port 443 is the standard port for HTTPS, which is essential for encrypting communication between web browsers and a secure website. (CompTIA Server+ Objectives SK0-004: 4.1)
Why other options are not essential:

25 (SMTP): Used for email transmission
53 (DNS): Used for domain name resolution
**3389 (RDP): ** Used for remote desktop connections
**8080 (Alternate HTTP): ** Sometimes used for web servers, but not the standard secure port

Question: 390

A systems administrator is provisioning a large number of virtual Linux machines that will be configured identically. The administrator would like to configure the machines quickly and easily but does not have access to an automation/orchestration platform. Additionally, the administrator would like to set up a system that can be used in the future, even on newer versions of the OS. Which of the following will best meet the administrator's requirements?

A. Deploying each server from a VM template
B. Using a kickstart file during installation
C. Configuring each server manually one at a time
D. Copying/pasting configuration commands into each server through SSH sessions
E. Configuring a single server and then creating clones of it

Answer: B

Explanation:

Kickstart Files (Linux): Kickstart files are configuration files that automate the Linux installation process. They contain pre-determined answers to installation prompts, allowing for identical and rapid deployment of multiple systems. (CompTIA Server+ Objectives SK0-004: 3.1, Red Hat documentation on Kickstart: https://access.redhat.com/documentation/)
Why other options are less ideal:
VM Template (A): Templates are useful for replicating the OS & some software, but might not capture all configurations.
Manual Configuration (C): Time-consuming and prone to errors when replicating across many servers.
Copy/Paste via SSH (D): Tedious, error-prone, and requires servers to be online before configuration. Cloning (E): Can work but has version compatibility risks if the OS of the cloned server isn't identical to the new ones.

Question: 391

A server located in an IDF of a paper mill reboots every other day at random times. Which of the following should the technician perform on the server first?

A. Check the power cables

B. Clean the fans.
C. Replace the RAM.
D. Reattach the CPU heat sink

Answer: A

Explanation:

In a situation where a server reboots randomly, the first step should be to check for any issues with the power supply. Random reboots can often be caused by intermittent power supply issues, which can be due to faulty power cables, loose connections, or problems with the power source itself. This is especially pertinent in environments like a paper mill where dust and debris might affect cable integrity. Since the issue occurs every other day and at random times, it's less likely to be caused by components that would typically fail due to overheating or other gradual issues (like RAM or CPU heat sink problems). Therefore, checking the power cables is the simplest and most direct first step to troubleshoot the issue.

Question: 392

An administrator receives an alert that one of the virtual servers has suddenly crashed. The administrator confirms the data center does not have any power failures and then connects to the remote console of the crashed server After connecting to the server console, which of the following should the administrator complete first?

A. Use the keyboard command AH+F12 to switch to the kernel log screen
B. Perform a hard reboot on the server and monitor the server startup
C. Collect a screenshot of the PSOD and note the details after the line detailing the OS version
D. Collect a core dump from the server and store locally before rebooting the hardware

Answer: C

Explanation:

When a virtual server crashes and presents a Purple Screen of Death (PSOD), the immediate response should be to document the incident. Collecting a screenshot of the PSOD is crucial as it

contains error codes and state information that can be used for diagnosing the root cause of the crash. Noting the details, especially those that come after the line detailing the OS version, can provide specific clues to what might have caused the server to crash. This is a standard best practice before rebooting the server, as it ensures that there is a record of the event to investigate and potentially prevent future occurrences. A hard reboot should only be done after this critical information has been recorded.

Question: 393

A technician is attempting to resolve an issue with a file server that is unable to download a file
Given the following output:

```
mt?35rver:-S la -2 Zvar/www/h.tinl/file
-rw-r—r— root root unconfined_u:object_r:sanba_5hare_t:s- /var wwn1html'file
```

Which of the following would best allow this file to be read?

A. chown
B. sestatus
C. setenforce
D. getenforce
E. chmod

Answer: E

Explanation:

The given output in the image indicates that the file is present, but the permissions may not allow it to be read. The output indicates '-rw - ', which means that the file is set to be readable and writable by the owner only, with no permissions for group or others. To allow the file to be read by users other than the owner, the file's permissions will need to be changed. Thechmod(change mode) command is used to change the file's permissions in Linux. For example,chmod 644 filewould change the permissions of the file to be readable by everyone and writable by the owner, which is typically what's required for a file server. It is always recommended to apply the least permissive settings that still allow the required operation to maintain security.

Question: 394

A technician replaces a single faulted disk in the following array RAID 10, Four 15K SAS HDD The technician replaces it from a disk in spare parts, and the array rebuilds the data in a few minutes. After the array rebuild is complete, the system reports the IOPS on the disk array have dropped by almost 60% Which of the following should the technician investigate first?

A. Check the RAID controller (or background rebuild tasks
B. Check the firmware version on the newly replaced disk.
C. Check the RPM speed on the newly replaced disk-
D. Check the cache settings on the RAID controller.

Answer: C

Explanation:

In RAID 10 arrays, disk performance is crucial, especially if they are high-speed 15K RPM SAS HDDs, as each disk in the array is part of a mirrored pair that also stripes data with another pair. When replacing a disk, it's essential that the new disk matches the specifications of the others, especially in terms of rotational speed (RPM). If the replaced disk is slower, it can significantly reduce the Input/Output operations per second (IOPS) of the entire array. This is because all disks need to work in tandem, and the slowest disk can become a bottleneck. Thus, checking the RPM of the newly replaced disk is a sensible first step to ensure it matches the performance of the other disks in the array.

Question: 395

A new company policy requires that any lost functionality must be restored within 24 hours in the event of a disaster. Which of the following describes this policy requirement?

A. MTBF
B. RTO
C. MTTR
D. RPO

Answer: B

Explanation:

Recovery Time Objective (RTO) refers to the target time set for the recovery of IT and business activities after a disaster has struck, which includes restoring server, network, and data access. The policy requirement mentioned in the question aligns with the definition of RTO, as it specifies the maximum allowable downtime or the time within which functionality must be restored. Mean Time Between Failures (MTBF) and Mean Time To Repair (MTTR) are metrics related to the reliability and

repair times of systems but do not specifically pertain to disaster recovery time frames. Recovery Point Objective (RPO) describes the acceptable amount of data loss measured in time, not the restoration of operations.

Question: 396

An administrator has been asked to copy files from a Windows server that may not conform to Windows file-naming standards. Which of the following would best facilitate the copy process?

A. Robocopy
B. SCP

C. Drag and drop
D. FTP

Answer: A

Explanation:

Robocopy (Robust File Copy) is a command-line tool in Windows that is designed for reliable copy or mirroring of files, and it can handle a broader range of file names and paths, including those that do not conform to traditional Windows file-naming standards. It's specifically designed to handle complex file copy demands and offers a wide range of options that can be tailored for different scenarios, which makes it suitable for the task mentioned. SCP (Secure Copy Protocol), Drag and Drop, and FTP (File Transfer Protocol) are all methods that can be used to copy files, but they might not handle non-standard Windows file names as well as Robocopy.

Question: 397

A technician is configuring a server rack that will hold ten blade servers. Which of the following safety concerns should be observed? (Select three).

A. Floor load limitations
B. Rack balancing
C. Proper lifting techniques
D. Power connector type
E. KVM placement
F. Cable management
G. UPS power requirements
H. PDU installation
I. Separate circuits for power

Answer: A,B,C

Explanation:

When configuring a server rack, it's important to consider:
A .Floor load limitations: Server racks can be extremely heavy, especially when filled with equipment like blade servers. It is crucial to ensure that the floor can handle the load to avoid structural damage or failure.
B .Rack balancing: Properly distributing the weight in a server rack is important for stability. Heavier equipment should generally be placed at the bottom to prevent the rack from becoming top-heavy and risking a tip-over.
C .Proper lifting techniques: Using correct lifting techniques when placing servers into a rack is vital to prevent personal

injury.

Question: 398

A virtual host has four NICs and eight VMs. Which of the following should the technician configure to enable uplink redundancy?

A. VM
B. vNIC
C. vSwitch
D. vCPU
E. vHBA

Answer: C

Explanation:

Uplink redundancy is a method used to ensure that if one physical network interface card (NIC) fails, the network connectivity for the virtual machines (VMs) does not go down. This is typically achieved by configuring multiple NICs to connect to a single virtual switch (vSwitch) and setting up NIC teaming or bonding. The vSwitch manages the internal network traffic between the VMs and the outside network by using the physical NICs assigned to it. By configuring the vSwitch with multiple NICs, you can create redundancy, so if one NIC fails, the other NICs can take over the traffic, ensuring continuous network connectivity.

Question: 399

An employee who was dismissed did not return company-issued equipment. Which of the following is the most important information the IT department needs to give to the legal department?

A. Labeling
B. Serial number
C. Warranty
D. Asset tag

Answer: D

Explanation:

The most important piece of information needed by the legal department in the event that an employee does not return company-issued equipment is the asset tag. The asset tag is a unique identifier that is used to track assets throughout their lifecycle. It allows the company to keep precise records of the assets, monitor their location, and manage their overall inventory. In legal situations, the asset tag can be used to prove ownership and aid in the recovery process of the equipment. The serial number is also important, but it is the asset tag that ties the equipment directly to the company's asset management system and is therefore the most crucial for the legal department. Warranty and labeling information are less critical from a legal perspective when it comes to unreturned equipment.

Question: 400

Which of the following authentication types defines credentials as "something you have"

A. Swipe pattern
B. PIN
C. Fingerprint
D. Smart card

Answer: D

Explanation:

The concept of authentication is rooted in the principle of verifying identity, which is commonly broken down into three categories: "something you know" (like a password or PIN), "something you have" (such as a smart card or a security token), and "something you are" (biometric data, for example, fingerprints). The question asks for the authentication type defined by "something you have."

Question: 401

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

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

Answer: A

Explanation:

Load balancing is a technique used to distribute workloads evenly across multiple servers, ensuring no single server is overwhelmed. This is especially important in environments where high availability and reliability are critical, such as when multiple servers are serving the same website. By doing so, load balancing improves the responsiveness and availability of applications or websites.

Load balancing refers to the process of distributing network or application traffic across multiple servers to ensure no single server becomes overwhelmed, thereby improving responsiveness and availability of applications or websites. In the scenario described, where four servers are set up to each serve the same website, the concept of load balancing is applied. This setup aims to distribute incoming requests evenly among the servers to maximize speed and capacity utilization while ensuring no one server is overburdened, which can lead to improved overall performance of the website. Options B, C, and D do not accurately describe the scenario of distributing traffic for the same website across multiple servers.

Question: 402

A new company policy requires that any data loss in excess of one hour is unacceptable in the event of a disaster. Which of the following concepts is being referred to in this policy?

A. MTTR
B. RTO
C. RPO
D. MTBF

Answer: C

Explanation:

The Recovery Point Objective (RPO) refers to the maximum tolerable period in which data might be lost from an IT service due to a major incident. The policy mentioned in the question highlights that data loss exceeding one hour is unacceptable, directly relating to the RPO concept. RPO is critical in disaster recovery and business continuity planning, indicating the age of the files that must be recovered from backup storage for normal operations to resume without significant losses. MTTR (Mean Time To Repair), RTO (Recovery Time Objective), and MTBF (Mean Time Between Failures) are related concepts but do not specifically address the amount of data loss that can be tolerated.

Question: 403

A server administrator is tasked with upgrading the network on a server to 40Gbps. After installing the card, which of the following connectors should the administrator use?

A. QSFP+
B. 10 GigE
C. SFP
D. SFP+

Answer: A

Explanation:

QSFP+ (Quad Small Form-factor Pluggable Plus) connectors are used in high-density, high-speed networking solutions such as 40 Gigabit Ethernet (40GbE) interfaces. When upgrading a server's network to 40Gbps, QSFP+ is the appropriate choice due to its capability to support such high-speed data transfer rates. 10 GigE, SFP, and SFP+ connectors are used for lower speed connections (10Gbps and below for SFP+ and 10 GigE, and even less for SFP), making them unsuitable for a 40Gbps network upgrade.

Question: 404

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

A. The media destruction method
B. The recycling policy
C. Asset management documentation
D. Document retention policy

Answer: C

Explanation:

Before decommissioning and wiping a file server, it's crucial to verify the asset management documentation. This documentation provides detailed records of the server's lifecycle, including procurement, usage, maintenance, and decommissioning information. Ensuring that asset management documentation is up-to-date and accurate is essential before proceeding with the server's decommissioning to maintain proper inventory control, comply with regulatory and organizational policies, and facilitate any potential audits. While the media destruction method, recycling policy, and document retention policy are important considerations in the decommissioning process, verifying asset management documentation is the first step to ensure the server is correctly identified and accounted for in the organization's asset registry.

Question: 405

An administrator is setting up a new employee's read/write access to a document on the file server. Currently, the user can open the file, but edits cannot be saved. Which of the following should the administrator do so the user can save the updated file while maintaining least-privilege access?

A. Give the user "read and execute" rights to the file.
B. Give the user "modify" rights to the file.
C. Give the user "list folder contents" rights to the folder.
D. Give the user "full control" rights to the folder.

Answer: B

Explanation:

To enable a user to both read and write to a file, the administrator needs to provide "modify" rights. This permission level allows the user not only to open and view the file (read) but also to make changes and save those changes back to the file (write). "Read and execute" rights would allow the user to open and run the file but not make any changes. "List folder contents" rights would enable the user to view the files within a folder but not make changes to them. "Full control" rights would provide the user with complete control over the file, including the ability to change permissions and delete the file, which exceeds the principle of least-privilege access necessary for this task.

Question: 406

An IT administrator is configuring ten new desktops without an operating system. The infrastructure contains an imaging server and operating system loaded on a USB, a DVD, and an SD card. Which of the following options would minimize the amount of time the administrator needs to load the operating system on every machine?

A. SD card B. Optical C. Network D. USB

Answer: C

Explanation:

Using a network-based deployment, such as network booting (PXE - Preboot Execution Environment) or imaging through a server, is the most efficient way to load operating systems onto multiple machines simultaneously. This approach minimizes the manual intervention required for each device, as the administrator can initiate the operating system installation or imaging process across all desktops at once through the network. In contrast, using an SDcard, DVD (Optical), or USB would require the administrator to physically move the media from one desktop to another, significantly increasing the setup time for each device.

Question: 407

A server administrator is tasked with resolving an issue with the server's local storage. The administrator turns on the server and only two out of the four drives are found. After several reboots, an additional hard drive connects and disconnects randomly. Which of the following is most likely the cause?

A. RAM
B. Power supply
C. Cooling failure
D. CPU

Answer: B

Explanation:

The server's local storage is experiencing problems. Only two out of four drives are found, and an additional hard drive connects and disconnects randomly.
Likely Cause: The most probable cause is a faulty power supply (Option B). Here's why:

When a power supply is failing or not providing sufficient power, it can lead to intermittent connectivity issues with drives.
Drives may not be recognized during boot or may disconnect randomly due to power fluctuations. RAM, CPU, and cooling failures typically result in different symptoms (e.g., blue screens, system instability, or overheating), which do not match the described issue.
Reference: CompTIA Server+ Guide, Chapter 3: Hardware, Section 3.2.1 (Power Supply Units)

Question: 408

A server administrator is replacing a faulty PSU. The management team has asked for a solution that prevents further downtime in the future. Which of the following can the server administrator implement?

A. Separate circuits
B. Load balancing
C. Server monitoring
D. Redundancy

Answer: D

Explanation:

A faulty PSU needs replacement, and the management team wants a solution to prevent future downtime.
Solution: Implement redundancy (Option D):
Redundant power supplies ensure that if one fails, the other takes over seamlessly.
Separate circuits, load balancing, and server monitoring are important but do not directly prevent PSU failures.
Reference: CompTIA Server+ Guide, Chapter 3: Hardware, Section 3.2.2 (Redundant Power Supplies)

Question: 409

A new virtual server was deployed in a perimeter network. Users have reported the time on the server has been incorrect. The engineer has verified the configuration, and the internal time servers are configured properly. Which of the following should the engineer do to resolve this issue?

A. Check the firewall rules.
B. Replace the CMOS battery in the server.
C. Restart the time servers.
D. Manually correct the time.

Answer: D

Explanation:

The virtual server's time is incorrect despite proper configuration of internal time servers. Resolution: The engineer should manually correct the time (Option D):
If the internal time servers are configured correctly, manual adjustment is necessary.
Checking firewall rules would help identify NTP (Network Time Protocol) issues.

Replacing the CMOS battery is unlikely to be the cause of incorrect time.
Reference: CompTIA Server+ Guide, Chapter 4: Networking, Section 4.2.2 (Time Synchronization)

Question: 410

A systems administrator is installing Windows on an 8TB drive and would like to create a single 8TB partition on the disk.
Which of the following options should the administrator use?

A. VMFS
B. ZFS
C. GPT
D. MBR
E. LVM

Answer: C

Explanation:

Scenario: Installing Windows on an 8TB drive, aiming for a single 8TB partition.
Solution: Use the GPT (GUID Partition Table) option (Option C):
GPT supports larger partitions (up to 18.4 million TB) compared to MBR (limited to 2TB).
VMFS, ZFS, and LVM are unrelated to the partitioning scheme in Windows.
Reference: CompTIA Server+ Guide, Chapter 3: Hardware, Section 3.3.1 (Partitioning)

Question: 411

Which of the following types of locks utilizes key fobs or key cards held against a sensor/reader to gain access?

A. Bolting door lock
B. Combination door lock
C. Electronic door lock
D. Biometric door lock

Answer: C

Explanation:

Lock Type: The type of lock utilizing key fobs or key cards is an electronic door lock (Option C).
Electronic locks use electronic credentials (such as key cards) for access control.
Reference: CompTIA Server+ Guide, Chapter 5: Security, Section 5.1.2 (Access Control)

Question: 412

A startup is migrating a stand-alone application that stores PII to the cloud. Which of the following should be encrypted?

A. Datain transit
B. Dataat rest
C. Databackups
D. Dataarchives

Answer: B

Explanation:

Scenario: Migrating a stand-alone application storing PII to the cloud.
Data to Encrypt: Data at rest (Option B) should be encrypted:
Data in transit (Option A) and backups (Option C) are also important but not the primary concern here.
Data archives (Option D) may or may not contain sensitive data.
Reference: CompTIA Server+ Guide, Chapter 5: Security, Section 5.3.1 (Data Encryption)

Question: 413

An administrator received an alert that a backup job has been unsuccessful in the previous three attempts. The administrator discovers the issue occurred while backing up a user's data on a network share. Which of the following actions would be best to allow the job to complete successfully?

A. Enabling open file backups in the backup job
B. Moving the user's data off the network share

C. Excluding the user's data from the backup
D. Changing the backup job to exclude certain file types

Answer: A

Explanation:

Enabling open file backups allows the backup software to copy files that are currently open or in use by applications. This ensures that even if a file is actively being modified during the backup process, it will still be included in the backup.

Advantages:

Ensures consistency of backed-up data, especially for files that are frequently accessed or modified.

Prevents data loss due to files being skipped during backup.

Reference:

CompTIA Server+ Guide, Chapter 6: Storage, Section 6.3.2 (Backup Types)

Microsoft Docs: Volume Shadow Copy Service (VSS)

Question: 414

An administrator is patching a file and print server After rebooting the server, it begins acting strangely. The administrator tries to open Print Management, but it will not open Upon inspection, the spooler service has been disabled. The administrator then notices the server has services that should be enabled but are disabled. Which of the following actions should the administrator take next to resolve the issue as quickly as possible?

A. Roll back the installed updates
B. Set the services to restart automatically and reboot the server
C. Enable the spooler service.
D. Reboot the server into safe mode

Answer: B

Explanation:

When a server is acting strangely after a patch and services that should be enabled are disabled, it's

often a good first step to set the services to restart automatically and then reboot the server. This can resolve many issues as it ensures that all services start correctly upon reboot

Question: 415

Which of the following will protect critical data during a natural disaster?

A. Off-site storage
B. Life-cycle management
C. Environmental controls
D. Data-at-rest encryption

Answer: A

Explanation:

Off-site storage involves storing data backups in a physically separate location from where the servers and primary data reside. This protects critical data in case the main site is damaged or destroyed due to a natural disaster, such as floods, fires, or earthquakes. Off-site storage can be in the form of cloud storage or physical tapes stored in a geographically distant location.

Off-site storage (Answer A):It ensures that a company's critical data remains accessible even if the local infrastructure is compromised. This is a common part of disaster recovery planning and is crucial for business continuity.

Life-cycle management (Option B):While important for managing data from creation to deletion, it does not specifically protect data during disasters.

Environmental controls (Option C):These manage physical conditions like temperature and humidity but do not address the risk of data loss during disasters.

Data-at-rest encryption (Option D):Encryption protects the confidentiality of stored data but does not ensure availability in the event of a disaster.

CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 4.1: Summarize disaster recovery methods and concepts.

Question: 416

A technician is preparing a deployment of servers to be used by staff at a remote location. Which of the following should the technician do to prevent access to the hardware configuration?

A. Enable an administrator account
B. Enable a UEFI password
C. Disable WOL
D. Enable encryption at rest

Answer: B

Explanation:

Enabling aUEFI (Unified Extensible Firmware Interface) passwordprevents unauthorized users from making changes to the

server's hardware configuration settings, such as boot order or device settings. This is crucial for protecting the integrity of the server at a remote location where physical security might be more difficult to enforce.

UEFI password (Answer B):It provides security at the firmware level, preventing changes to low-level configurations unless the correct password is provided.

Administrator account (Option A):While important for OS-level access, it doesn't prevent someone with physical access from altering hardware settings via UEFI/BIOS.

Disabling WOL (Option C):Wake-on-LAN (WOL) allows a device to be powered on remotely. Disabling it can help with security but does not prevent hardware configuration changes.

Encryption at rest (Option D):Encryption protects data on the server but does not prevent hardware configuration access.

CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 2.1: Install and configure server operating systems.

Question: 417

An administrator is receiving reports that users in a remote office are unable to log in to the domain. The network appears to be up between the sites, which are in different states. Which of the following is the most likely cause of the issue?

A. The user accounts are locked out
B. The NTP on the client side is misconfigured
C. MFA is not working
D. The wrong time zone was set in the remote office

Answer: B

Explanation:

A misconfiguredNTP (Network Time Protocol)on the client side can cause authentication issues because Active Directory relies on accurate time synchronization between clients and domain controllers for security purposes. If the clocks are out of sync, even by a few minutes, it can prevent users from authenticating to the domain.

NTP misconfiguration (Answer B):If the client machines have incorrect time settings, it can disrupt Kerberos-based authentication, which is time-sensitive.

Locked user accounts (Option A):While it can prevent login, it would affect individual users rather

than a whole office.

MFA (Option C):Multi-factor authentication not working would typically cause a login issue but is less likely across an entire remote office.

Wrong time zone (Option D):While incorrect time zones can cause confusion, they usually don't prevent logins as long as the NTP server synchronizes the actual time correctly.

CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 2.4: Troubleshoot common server problems.

Question: 418

A user is unable to access the database server from a workstation. All other workstations are able to access the database server. The technician issues the following command to check the workstation:

userGhost:!# nmap localhost
Nmap scan report for localhost (127.0.0.1)
Host is up (0.001 ms latency).
Not shown: 1023 closed ports
PORT STATE SERVICE
3306 open mysql
Which of the following should the technician perform next to allow access?

A. Check the database server
B. Change the database password
C. Open port 3306 on the server
D. Edit the network firewall rules

Answer: D

Explanation:

Sinceport 3306(MySQL) is open on the local workstation as confirmed by the Nmap scan, the next logical step is to checknetwork firewall rules. It's possible that the workstation's firewall or networklevel security rules are blocking outbound or inbound traffic on port 3306, which is required for database access.
Edit the network firewall rules (Answer D):Ensuring the firewall allows traffic on port 3306 is the appropriate next step, especially since other workstations are not having the same issue.
Checking the database server (Option A):This is unnecessary because other workstations can access the server, confirming the server itself is functioning correctly.
Changing the database password (Option B):This would not resolve the issue of one workstation being unable to connect while others can.
Opening port 3306 on the server (Option C):The problem isn't server-side since other machines can connect without issue.
CompTIA Server+ Reference:This topic relates toSK0-005 Objective 2.4: Troubleshoot common server problems.

Question: 419

A technician notices that every time a server is powered on, it turns off after several minutes. After reviewing logs, the technician notices the server registers execution of the shutdown. Which of the following should the technician do to fix this issue?

A. Reset the memory modules
B. Check for fan failure

C. Change the VM password
D. Set credentials for the remote console

Answer: B

Explanation:

The server is shutting down after being powered on, which could indicateoverheating. This is a common issue whenfans fail, causing the CPU or other components to overheat and forcing the system to shut down to protect itself from damage.

Check for fan failure (Answer B):Ensuring the server's cooling system is functioning properly is crucial. Overheating due to a fan failure can cause the system to shut down automatically.

Resetting memory modules (Option A):While memory issues can cause system instability, they generally do not lead to immediate shutdowns as described.

Changing the VM password (Option C):This is unrelated to the shutdown issue.

Setting credentials for the remote console (Option D):This is irrelevant to the described problem of server shutdown.

CompTIA Server+ Reference:This topic is related toSK0-005 Objective 3.3: Diagnose hardware and software issues.

Question: 420

A server administrator needs to set up active-passive load balancing for an environment with multiple network paths. The traffic should be directed to one path and switched only in the event that the original path becomes unavailable. Which of the following describes this scenario?

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

Answer: B

Explanation:

In anactive-passive load balancingconfiguration, one path (or server) handles all the traffic while the other remains in standby mode, ready to take over if the active path becomes unavailable. This setup is typically monitored via aheartbeatmechanism, where the standby server checks the availability of the active server and takes over when it detects a failure.

Heartbeat (Answer B):This term refers to the monitoring process that checks the availability of the active path in an active-passive configuration.

Most recently used (Option A):This describes a load-balancing method where the most recently used path or server is prioritized, but it doesn't describe an active-passive setup.

Link aggregation (Option C):This refers to combining multiple network connections for increased throughput, not for failover purposes.

Round robin (Option D):This load-balancing method alternates traffic between all available paths, which doesn't match the active-passive description.

CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 1.4: Summarize methods used to manage network connections.

Question: 421

A junior administrator reported that the website used for anti-malware updates is not working. The senior administrator then discovered all requests to the anti-malware site are being redirected to a malicious site. Which of the following tools should the senior administrator check first to identify the potential cause of the issue?

A. Data loss prevention
B. File integrity monitor
C. Port scanner
D. Sniffer

Answer: D

Explanation:

Asniffer(also known as a packet analyzer) is a tool that captures and inspects data packets traveling across the network. In this case, using a sniffer would help identify suspicious or malicious redirection of traffic, possibly caused by a man-in-the-middle attack, DNS hijacking, or malware. Sniffer (Answer D):This tool will allow the senior administrator to inspect the network traffic and identify if and how requests to the anti-malware website are being intercepted or redirected. Data loss prevention (Option A):DLP tools focus on preventing data leakage rather than analyzing traffic redirection.
File integrity monitor (Option B):This checks for unauthorized changes to files, which may not directly address network traffic redirection.

Port scanner (Option C):A port scanner would only identify open ports on devices, which is unrelated to the redirection issue.
CompTIA Server+ Reference:This topic is addressed underSK0-005 Objective 4.2: Explain server roles and their purposes.

Question: 422

Which of the following is a benefit of failover NIC teaming?

A. Decreased latency
B. Increased transmission size
C. Increased system resources
D. Increased fault tolerance

Answer: D

Explanation:

Failover NIC teaming(Network Interface Card teaming) is a method used to providefaulttoleranceby combining two or more NICs into a single logical interface. If one NIC fails, the other NIC(s) take over, ensuring continuous network connectivity.

Increased fault tolerance (Answer D):The primary benefit of failover NIC teaming is to ensure that network communication is not interrupted when one NIC fails. The remaining NICs take over to maintain uptime.

Decreased latency (Option A):NIC teaming is not primarily used to reduce latency; it's focused on redundancy and fault tolerance.

Increased transmission size (Option B):NIC teaming does not directly increase the size of transmissions.

Increased system resources (Option C):NIC teaming doesn't increase system resources like CPU or memory, but it ensures network redundancy.

CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 1.4: Summarize methods used to manage network connections.

Question: 423

A systems administrator notices that a SAN is running out of space. There is no additional funding in the budget to upgrade the storage space. Which of the following will significantly reduce the storage space with the least effort? (Select two).

A. Configuring compression
B. Deleting old and unused files
C. Increasing the number of IOPS
D. Decreasing the RAID level
E. Enabling deduplication
F. Upgrading the filesystem type

Answer: A,E

Explanation:

To free up storage space on a SAN with minimal effort and no additional cost, the following methods can be used:

Configuring compression (Answer A):Compression reduces the size of files stored on the SAN, allowing more data to be stored in the same space.

Enabling deduplication (Answer E):Deduplication identifies and removes duplicate copies of data, reducing the amount of space required for storage.

Deleting old and unused files (Option B):While effective, this requires manual effort and may not be the best long-term solution.

Increasing the number of IOPS (Option C):This relates to improving performance, not reducing storage space.

Decreasing the RAID level (Option D):Lowering the RAID level could reduce redundancy, which may free up space but

also decreases fault tolerance.

Upgrading the filesystem type (Option F):This is a complex process that doesn't directly reduce storage usage.

CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 3.2: Summarize storage technologies and solutions.

Question: 424

A college is planning for disaster recovery and needs to have access at all times to student data, which contains PII (Personally Identifiable Information). Which of the following would be the most appropriate for the college?

A. A warm site with private cloud backup
B. A warm site with public cloud backup
C. A cold site with public cloud backup
D. A hot site with private cloud backup

Answer: D

Explanation:

Ahot site with private cloud backupis the most suitable for environments that requirecontinuous accessto critical data, like student data containingPII. A hot site is fully operational and ready to take over immediately in case of a disaster. Using aprivate cloudprovides more control and security, which is important when handling sensitive data like PII.

Hot site with private cloud backup (Answer D):This ensures minimal downtime and robust security, which is critical for handling PII.

Warm site (Options A & B):A warm site has some infrastructure ready but requires some setup before becoming operational, which introduces some delay.

Cold site (Option C):A cold site has the least preparation and requires significant time to become operational, which would not meet the requirement for constant availability.

CompTIA Server+ Reference:This topic is addressed inSK0-005 Objective 4.1: Summarize disaster recovery methods and concepts.

Question: 425

Which of the following should be created to understand how long data is stored and how frequently data backups should be scheduled?

A. Retention policies
B. Service-level agreement
C. Life-cycle management
D. Mean time to recover

Answer: A

Explanation:

Retention policiesdefine how long data should be kept before being archived or deleted and help in determining the frequency of backups to ensure data is available when needed. These policies are crucial for compliance, storage management, and disaster recovery planning.

Retention policies (Answer A):These directly address how long data is kept and how often it should be backed up.

Service-level agreement (Option B):SLAs define performance expectations and uptime, not specific data retention or backup schedules.

Life-cycle management (Option C):This refers to managing data from creation to disposal but doesn't specifically address backup frequency or retention periods.

Mean time to recover (Option D):This metric measures how long it takes to restore services after a failure, not backup scheduling.

CompTIA Server+ Reference:This topic relates toSK0-005 Objective 4.2: Explain the importance of data security concepts.

Question: 426

While a technician is troubleshooting a performance issue on a database server, users are disconnected from the database. An administrator is asked to intervene and restore access. Which of the following steps should the administrator take first?

A. Revert any patches or updates on the server from the past 24 hours
B. Reproduce the issue in a test environment and confirm the database fails in the same manner
C. Perform a root cause analysis and report the issue to management
D. Collect all logs from the system and review the actions the technician performed
E. Perform a quick backup of the system to prevent any further issues

Answer: D

Explanation:

The first step when addressing an urgent issue like a database disconnection should be to gather information to understand what actions caused the problem.Collecting all logs and reviewing the technician's actions (Answer D)will help identify what went wrong and guide appropriate remediation. Restoring access quickly is the priority, and without reviewing logs, it would be hard to identify the root cause or prevent further damage.

Reverting patches (Option A):Reverting patches might cause more issues if the problem isn't directly related to a recent update.

Reproducing the issue (Option B):Reproduction is useful in a controlled environment but not practical as an immediate first step in production.

Root cause analysis (Option C):This is necessary but should come after restoring access.

Performing a quick backup (Option E):While backups are crucial, this step might not help resolve the current access problem and can delay troubleshooting.

CompTIA Server+ Reference:This topic relates toSK0-005 Objective 3.4: Explain troubleshooting theory and methodologies.

Question: 427

A technician needs to replace two RAID controllers on a database server as part of an upgrade. The server has six external storage arrays and eight internal disks that are controlled by the two RAID controllers. The technician completes the replacement and powers the systems back on, but the server OS detects several missing disks in the configuration. Which of the following steps should the technician take first to resolve this issue?

A. Reboot to the RAID controller BIOS and rescan for attached disks and arrays
B. Start the motherboard UEFI and confirm the RAID controller start order is correct
C. Turn off the external arrays and reboot the server to redetect the internal disks first
D. Turn off all components and confirm all external and internal cables before rebooting

Answer: A

Explanation:

When new RAID controllers are installed and disks go missing, the first logical step is toreboot to the RAID controller BIOS and rescan for attached disks and arrays (Answer A). This will allow the RAID controller to detect the connected drives and rebuild the RAID configuration as expected.
Rebooting to RAID controller BIOS (Answer A):This step ensures that all the attached disks and arrays are properly detected and mapped to the RAID configuration.
Checking UEFI settings (Option B):This might be useful later, but initially, the RAID controller BIOS should be checked first to verify disk detection.
Turning off external arrays (Option C):While this could be helpful in some scenarios, it's not the most efficient first step.
Checking cables (Option D):This could be necessary, but since the controllers were just replaced, software/BIOS configurations should be checked first.
CompTIA Server+ Reference:This topic relates toSK0-005 Objective 3.1: Install and configure server components.

Question: 428

A systems engineer is configuring a new VLAN for expansion of a campus network. The engineer configures a new DHCP scope on the existing Windows DHCP server cluster and activates the scope for the clients. However, new clients in the area report they are not receiving any DHCP address information. Which of the following should the engineer do first?

A. Confirm the client PCs are using the correct speed on the NIC
B. Confirm the network uses the same NTP server as the clients
C. Confirm the network has a helper address for the DHCP server
D. Confirm the DHCP server is authorized in Active Directory

Answer: C

Explanation:

When clients on a different VLAN cannot receive DHCP addresses, it's often because there is noDHCP helper addressconfigured on the router or switch. ADHCP relay/helper addressforwards DHCP requests from clients on the VLAN to the DHCP server, which may reside on a different network. Confirming the network has a helper address (Answer C):This should be the first step to ensure that DHCP requests from the VLAN are correctly routed to the DHCP server.

Checking client NIC speed (Option A):While important, incorrect NIC speed wouldn't affect the ability to obtain an IP address from DHCP.

Checking NTP server (Option B):Time synchronization is unrelated to DHCP issues.

Authorizing the DHCP server (Option D):This is essential for DHCP operation, but since this is an existing DHCP cluster, the helper address is more likely the issue.

CompTIA Server+ Reference:This topic is related toSK0-005 Objective 1.2: Manage server network connections.

Question: 429

A technician is setting up a repurposed server. The minimum requirements are 2TB while ensuring the highest performance and providing support for one drive failure. The technician has the following six drives available:

500GB, 10,000rpm
600GB, 10,000rpm
500GB, 7,200rpm
500GB, 10,000rpm
600GB, 15,000rpm
600GB, 10,000rpm

Which of the following drive selections should the technician utilize to best accomplish this goal?

A. 1, 2, 4, and 6
B. 1, 2, 3, 5, and 6
C. 1, 2, 4, 5, and 6
D. 1, 2, 3, 4, and 6

Answer: C

Explanation:

The goal is to achieve2TB total storage,high performance, andsupport for one drive failure. Using RAID 5 or RAID 10 can provide redundancy while maximizing performance.

1, 2, 4, 5, and 6 (Answer C):This selection includes the highest-capacity and highest-performance drives (600GB and 500GB at high RPM speeds). Using these drives will provide the necessary capacity (500GB + 600GB + 500GB + 600GB + 600GB = 2.8TB) and performance with one drive redundancy (RAID 5 or 10).

Other options (A, B, D):These combinations either lack the necessary capacity or mix slower drives (7,200rpm), reducing overall performance.

CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 1.5: Explain storage technologies and RAID configurations.

Question: 430

Which of the following should a server administrator use when writing a script with a function that needs to be run ten times?

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

Answer: A

Explanation:

Aloopis a programming construct used to repeat a block of code multiple times. In this case, if a function needs to run ten times, a loop (such as afor looporwhile loop) would be the appropriate choice.
Loop (Answer A):This allows the function to be executed repeatedly without writing redundant code.
Variable (Option B):A variable is used to store data, but it doesn't handle repetition.
Comparator (Option C):Comparators are used to compare values, not for repeating code.
Conditional (Option D):Conditionals (if, else) are used for decision-making, not repeating actions multiple times.
CompTIA Server+ Reference:This topic is related toSK0-005 Objective 1.1: Understand basic scripting and automation concepts.

Question: 431

An administrator notices a server is ofline. Upon checking the console, the administrator discovers the server is stuck at:

Configuring Memory .....
After a reboot, the server still exhibits the same behavior. The administrator is able to log in to the OOB remote management but is unable to log in to the server. Which of the following is the most likely cause of this issue?

A. A DIMM has failed
B. The VRAM is insufficient
C. The RAID cache has failed
D. The BIOS needs to be updated

Answer: A

Explanation:

The error message "Configuring Memory..." and the fact that the server is stuck at this point strongly suggest thata DIMM

(memory module) has failed. Memory issues are commonly diagnosed when servers hang during the memory configuration phase of boot.

DIMM failure (Answer A):When a memory module fails, the server can get stuck at the memory initialization process during boot.

VRAM insufficiency (Option B):VRAM (video memory) issues would typically cause display-related errors, not memory configuration issues.

RAID cache failure (Option C):RAID cache issues would lead to storage errors but not prevent the system from configuring memory.

BIOS update (Option D):While BIOS issues can cause boot failures, the error message suggests a memory-specific issue.

CompTIA Server+ Reference:This topic relates toSK0-005 Objective 3.3: Diagnose hardware problems.

Question: 432

A technician is configuring a new server with four disks for the development team. The requirements are disk redundancy and maximum usable disk capacity. Which of the following RAID levels should be used for this server?

A. 0
B. 1
C. 5
D. 10

Answer: C

Explanation:

RAID 5offers bothdisk redundancyandmaximum usable disk capacityby using block-level striping with parity. RAID 5 allows one drive to fail while still retaining data, and it uses only one disk’s worth of space for parity, maximizing the usable capacity.

RAID 5 (Answer C):Provides redundancy with only one disk's space used for parity, maximizing storage.

RAID 0 (Option A):Provides no redundancy; if one drive fails, all data is lost.

RAID 1 (Option B):Mirrors data, so redundancy is provided, but only half the capacity is usable.

RAID 10 (Option D):Combines RAID 1 and 0 for high performance and redundancy but requires half the disk space for mirroring.

CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 1.5: Explain RAID levels and their benefits.

Question: 433

A server administrator created a new script and included the path to the script binary as the first line of the script. Which of the following scripting languages did the administrator most likely use?

A. Batch
B. Java
C. Bash
D. PowerShell

Answer: C

Explanation:

InBash(a Unix/Linux shell scripting language), it is common to include theshebang (#!) followed by the path to the interpreter(e.g., #!/bin/bash) on the first line of the script. This tells the operating system which interpreter to use to execute the script.

Bash (Answer C):Bash scripts often start with a shebang line that specifies the path to the binary (#!/bin/bash).
Batch (Option A):Batch scripts (used in Windows) do not require a path to the interpreter on the first line.
Java (Option B):Java is a compiled language, not a scripting language, so this does not apply.
PowerShell (Option D):PowerShell scripts (.ps1 files) do not typically require specifying the interpreter path on the first line.
CompTIA Server+ Reference:This topic is related toSK0-005 Objective 1.1: Understand scripting basics and the use of interpreters.

Question: 434

Which of the following describes the concept of allocating more resources than what is available on a hypervisor?

A. Direct access
B. Overprovisioning
C. Link aggregation
D. Component redundancy
E. Scalability

Answer: B

Explanation:

Overprovisioningrefers to allocating more resources (e.g., CPU, memory) to virtual machines on a hypervisor than the physical host actually has available. This is done because not all virtual machines are likely to use their maximum resources simultaneously. Overprovisioning helps optimize resource utilization on a hypervisor but must be managed carefully to avoid performance degradation. Overprovisioning (Answer B):It allows for more efficient use of resources, but overcommitting too much can lead to performance issues.
Direct access (Option A):Refers to direct hardware access, which is unrelated to resource allocation. Link aggregation (Option C):Refers to combining multiple network connections, unrelated to

hypervisor resources.

Component redundancy (Option D):Refers to using backup components for fault tolerance, not related to overprovisioning.

Scalability (Option E):Refers to the ability to increase or decrease resources as needed but doesn't describe overcommitting resources.

CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 1.3: Explain virtualization concepts and benefits.

Question: 435

Which of the following backup methods can be performed while a server is running, will not interrupt files in use, and can be used to fully restore the server if needed?

A. Snapshot
B. Archive
C. Open file
D. Differential

Answer: A

Explanation:

Asnapshotis a point-in-time image of a system's state that can be created while the system is running, without interrupting operations. Snapshots can be used to restore the entire system to the exact state at the time of the snapshot. They are commonly used in virtualized environments and allow administrators to roll back to a previous state if needed.

Snapshot (Answer A):Captures the entire system's state and can be taken without downtime.

Archive (Option B):Refers to long-term storage, not necessarily something that supports live backups. Open file (Option C):Refers to a method that handles open files during backup but doesn’t provide full system restoration capabilities.

Differential (Option D):Backs up only the data changed since the last full backup but still requires a full backup for complete system restoration.

CompTIA Server+ Reference:This topic is related toSK0-005 Objective 4.1: Explain backup and recovery methods.

Question: 436

A server administrator notices drive C on a critical server does not have any available space, but plenty of unallocated space is left on the disk. Which of the following should the administrator perform to address the availability issue while minimizing downtime?

A. Decrease the size of the page/swap file
B. Reformat drive C to use all the disk space
C. Enable whole disk encryption on the hard drive

D. Enforce disk quotas on shares within that server
E. Partition the unallocated space to provide more available space

Answer: E

Explanation:

If drive C is running out of space and there isunallocated spaceavailable on the disk, the best solution is topartition the unallocated spaceand extend drive C to include the new partition. This can be done without significant downtime and will immediately provide more space for the system.

Partition the unallocated space (Answer E):This is the simplest and least disruptive method to increase the available space on drive C.

Decreasing the page/swap file (Option A):This may free up some space but can negatively impact system performance.

Reformatting (Option B):Reformatting would result in significant downtime and data loss, making it an impractical option.

Whole disk encryption (Option C):This does not solve the issue of disk space.

Enforcing disk quotas (Option D):Disk quotas manage space usage but won't free up space on drive C. CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 3.3: Troubleshoot common server problems.

Question: 437

Which of the following types of physical security controls would most likely be a target of a social engineering attack?

A. A security guard
B. An access control vestibule
C. Perimeter fencing
D. Biometric locks
E. Bollards

Answer: A

Explanation:

Asecurity guardis a human element in physical security, making them susceptible tosocial engineering attacks. Social engineering exploits human behavior, and a guard can be tricked into allowing unauthorized access through persuasion, manipulation, or deception.

Security guard (Answer A):Human elements are the most vulnerable to social engineering techniques like impersonation or manipulation.

Access control vestibule (Option B):This is a physical security barrier, which is harder to exploit through social engineering.

Perimeter fencing (Option C):This is a static physical barrier, not susceptible to social engineering.

Biometric locks (Option D):These rely on biological data and are not susceptible to social engineering in the same way a human would be.

Bollards (Option E):Physical barriers that are not vulnerable to social engineering.

CompTIA Server+ Reference:This topic relates toSK0-005 Objective 4.4: Implement physical security controls.

Question: 438

Which of the following factors would most likely impact the selection of an organization's cloud provider?

A. Industry standards
B. Government regulations
C. Company policy
D. Organizational procedures

Answer: B

Explanation:

Government regulationsoften dictate the legal requirements for data storage, privacy, and security, which can greatly impact the selection of a cloud provider. Compliance with regulations like GDPR, HIPAA, or other local laws is critical when choosing a cloud provider to avoid legal repercussions. Government regulations (Answer B):Cloud providers must comply with legal and regulatory requirements, making this a significant factor in provider selection.

Industry standards (Option A):While important, standards can be more flexible and are often not legally binding.

Company policy (Option C):Internal policies are important but usually stem from the need to comply with regulations.

Organizational procedures (Option D):Procedures help guide operations but don't typically dictate the choice of cloud providers.

CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 1.6: Explain the importance of cloud-based concepts and services.

Question: 439

A bad actor leaves a USB drive with malicious code on it in a company's parking lot. Which of the following describes this scenario?

A. Hacking
B. Insider threat
C. Phishing
D. Social engineering

Answer: D

Explanation:

Leaving a USB drive in a company parking lot is a classic example ofsocial engineering, where the bad actor relies on human curiosity to prompt someone to pick up the USB drive and insert it into their computer, potentially infecting the system with malicious code.

Social engineering (Answer D):The attacker manipulates human behavior (in this case, curiosity) to exploit security weaknesses.

Hacking (Option A):Hacking refers to directly breaching security systems or exploiting software vulnerabilities, not manipulating humans.

Insider threat (Option B):This involves a legitimate insider within the organization carrying out malicious activities, which isn't the case here.

Phishing (Option C):Phishing typically involves emails or messages designed to deceive individuals into providing sensitive information.

CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 4.2: Explain server security concepts and best practices.

Question: 440

Which of the following describes when a site is considered a warm site?

A. It has basic technical facilities connected to it.
B. It has faulty air conditioning that is awaiting service.
C. It is almost ready to take over all operations from the primary site.
D. It is fully operational and continuously providing services.

Answer: C

Explanation:

Awarm siteis a backup location that haspartial infrastructureready to support operations in the event of a failure at the primary site. It isalmost ready to take overbut requires some configuration or installation of data backups and software before it can become fully operational.

Warm site (Answer C):This site has the necessary equipment but requires some setup, making it a middle ground between cold and hot sites.

Basic technical facilities (Option A):This more accurately describes a cold site.

Faulty air conditioning (Option B):This is irrelevant to the definition of a warm site.

Fully operational (Option D):This describes ahot site, which is continuously ready to provide full services.

CompTIA Server+ Reference:This topic is related toSK0-005 Objective 4.1: Summarize disaster recovery methods and concepts.

Question: 441

Which of the following supports virtualization?

A. Type 1 hypervisor
B. Bare-metal installation
C. Server-level redundancy
D. Active-active load balancing

Answer: A

Explanation:

AType 1 hypervisor(also known as a bare-metal hypervisor) is installed directly on the hardware and is designed to manage multiple virtual machines. It is the most common type of hypervisor used in virtualization environments.

Type 1 hypervisor (Answer A):This is a core technology that enables virtualization by running virtual machines directly on the physical hardware.

Bare-metal installation (Option B):Refers to installing an OS or hypervisor directly on hardware, but doesn't inherently support virtualization unless it's a hypervisor.

Server-level redundancy (Option C):This provides fault tolerance but is not related to virtualization. Active-active load balancing (Option D):Refers to distributing workloads across multiple servers, not directly tied to virtualization.

CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 1.3: Explain virtualization concepts.

Question: 442

The management team has requested that new software licenses be purchased out of the capital budget as one-time, non-renewing expenses this year. Which of the following types of software licenses would most likely be used to meet this request?

A. Open-source
B. Subscription
C. Volume
D. Perpetual

Answer: D

Explanation:

Aperpetual licenseis a one-time purchase that allows the software to be used indefinitely without recurring costs. This meets the requirement of a non-renewing, one-time capital expense.

Perpetual (Answer D):This is a one-time purchase and does not require renewal, making it ideal for capital expenditure.

Open-source (Option A):Open-source software is usually free but may not meet specific licensing requirements.

Subscription (Option B):This typically involves recurring payments, which contradicts the nonrenewing expense requirement.

Volume (Option C):Volume licenses are used for purchasing multiple licenses at a discount but may still involve subscription-based renewals.

CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 2.1: Explain licensing concepts for software.

Question: 443

After installing an OS on a new server, the administrator realizes the server does not have any network connectivity. The administrator checks the network cable, and it seems to be transferring dat
a. Which of the following should be checked next?

A. NIC firmware
B. IPv4 options
C. DHCP options
D. Firewall options
E. DNS settings

Answer: B

Explanation:

If the network cable is functioning but there is no network connectivity, the next step would be to check theIPv4 options. This includes verifying that the server has a correct IP address, subnet mask, and gateway configuration.
IPv4 options (Answer B):Incorrect IP configuration is a common reason for network connectivity issues.
NIC firmware (Option A):Firmware issues are less likely if the NIC is detecting the cable and appears operational.
DHCP options (Option C):DHCP could be checked if automatic IP assignment is being used, but verifying IPv4 options manually is the first step.
Firewall options (Option D):Firewalls can block connectivity but should be checked after confirming IP settings.
DNS settings (Option E):DNS issues usually affect domain name resolution, not basic network connectivity.
CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 3.1: Install and configure server components.

Question: 444

A technician wants to duplicate a physical server to a remote private cloud for disaster recovery purposes. Which of the following techniques would best accomplish this goal?

A. V2V
B. P2V
C. V2P
D. P2P

Answer: B

Explanation:

P2V (Physical to Virtual)refers to converting a physical server into a virtual machine, which can then be deployed to a remote cloud environment. This is the best method for duplicating a physical server to a virtual environment for disaster recovery.

P2V (Answer B):This is the process of converting a physical server to a virtual machine for migration to the cloud.

V2V (Option A):Refers to virtual-to-virtual migration, which is not applicable in this case.

V2P (Option C):Refers to converting a virtual machine back to physical, which is not needed here.

P2P (Option D):Physical-to-physical migration is not relevant for cloud-based disaster recovery.

CompTIA Server+ Reference:This topic relates toSK0-005 Objective 1.3: Explain virtualization concepts.

Question: 445

Which of the following would allow a server administrator to ensure all maximum available resources are being utilized?

A. Overprovisioning
B. Scalability
C. Thin clients
D. Resource Monitor

Answer: D

Explanation:

Resource Monitoris a tool that allows administrators to monitor the system's CPU, memory, disk, and network usage in real-time, ensuring that maximum resources are being efficiently utilized.

Resource Monitor (Answer D):This tool provides real-time insights into resource utilization and can help ensure resources are not under- or over-utilized.

Overprovisioning (Option A):Refers to allocating more resources than physically available but doesn't directly monitor resource usage.

Scalability (Option B):Refers to the ability to increase or decrease resources based on demand.

Thin clients (Option C):Refer to lightweight computers that depend on servers for processing power, unrelated to resource utilization monitoring.

CompTIA Server+ Reference:This topic is related toSK0-005 Objective 2.4: Monitor server performance.

Question: 446

Which of the following documents would explain the consequences of server downtime?

A. Service-level agreement
B. Business continuity plan
C. Disaster recovery plan
D. Business impact analysis

Answer: D

Explanation:

Abusiness impact analysis (BIA)is a document that outlines the potential effects of downtime on business operations. It identifies critical business functions and estimates the impact of disruption on the organization's ability to function.

Business impact analysis (Answer D):This document is specifically designed to assess the consequences of downtime and other disruptions.

Service-level agreement (Option A):Defines the expected level of service between a provider and a client but doesn't directly explain downtime consequences.

Business continuity plan (Option B):Outlines how the business will continue operating during a disruption but doesn't focus solely on the consequences of downtime.

Disaster recovery plan (Option C):Focuses on restoring systems after a disaster but doesn't outline the specific impact of downtime.

CompTIA Server+ Reference:This topic relates toSK0-005 Objective 4.1: Explain disaster recovery concepts.

Question: 447

Several mobile users are reporting intermittent performance issues when attempting to access network shares on the file server. After some investigation, the server administrator notices the server resources are running at maximum capacity, even during non-peak usage times. A recent port scan of the network identified this server as having too many unnecessary ports open to the public. Which of the following is the most likely cause of the performance issues?

A. Incorrect share permissions for mobile users
B. Improper privilege escalation
C. A virus infection
D. A misconfigured firewall rule

Answer: C

Explanation:

The fact that the server is running at maximum capacity, combined with the open ports, suggests that the server may have been compromised bya virus infection. Malware can exploit unnecessary open ports, leading to resource exhaustion and performance issues.

Virus infection (Answer C):Malware often exploits open ports and consumes system resources, leading to performance degradation.

Incorrect share permissions (Option A):This would affect access but not resource utilization.

Improper privilege escalation (Option B):This refers to unauthorized users gaining higher access rights, but it doesn't explain the high resource usage.

Misconfigured firewall rule (Option D):A misconfigured firewall could allow unnecessary traffic, but it's less likely to cause performance issues unless combined with malware.

CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 4.2: Explain server security concepts and best practices.

Question: 448

Multiple users have reported an issue accessing files on a specific server. Which of the following should be the first step in troubleshooting this issue?

A. Notify the key stakeholders of an outage.
B. Immediately execute a backup of the server.
C. Identify changes made to the environment.
D. Correct the permissions on the folder containing the files.

Answer: C

Explanation:

Comprehensive Detailed Explanation with all CompTIA Server+ SK0-005 Study Guide Reference:When troubleshooting server access issues, identifying recent changes to the environment is critical. Changes such as updates, patches, configuration modifications, or network adjustments could be the root cause of the issue. This aligns with the best practices of the ITIL framework and troubleshooting methodologies outlined in the CompTIA Server+ study materials: Step 1: Gather information and identify the scope of the problem.
Step 2: Look for recent changes to identify potential triggers.This approach ensures that corrective actions are based on an understanding of what caused the problem, avoiding unnecessary disruptions or incorrect fixes.

Question: 449

Under which of the following should a technician implement scripting?

A. Authenticating users
B. Change management
C. Business impact analysis
D. Step-by-step tasks

Answer: D

Explanation:

Comprehensive Detailed Explanation with all CompTIA Server+ SK0-005 Study Guide Reference:Scripting is commonly used to automate repetitive step-by-step tasks in server management, such as:
Automating system configurations.
Running maintenance jobs like backups or log file cleanup.
Performing administrative tasks like provisioning new users or setting permissions.
This automation improves efficiency and reduces the likelihood of human error. Referencing the SK0- 005 guide, scripting is a best practice for simplifying tasks that involve consistent, repeatable processes.

Question: 450

An administrator connects a server to an external gigabit switch, but the server can only get a speed of 100Mbps. Which of the following is the most likely cause of this issue?

A. Defective network cable
B. Port security
C. Incorrect IP subnet
D. Wrong VLAN

Answer: A

Explanation:

Comprehensive Detailed Explanation with all CompTIA Server+ SK0-005 Study Guide Reference:The most likely cause of a speed mismatch (100Mbps instead of 1Gbps) is a defective or incompatible network cable. Gigabit Ethernet requires at least Cat5e or Cat6 cables; using older cables (e.g., Cat5) or damaged cables can limit speeds to 100Mbps. Port securityandVLAN misconfigurationwould typically prevent communication entirely, not reduce speed. Incorrect IP subnetdoes not affect link speed.Checking and replacing the network cable is a logical first troubleshooting step in this scenario.

Question: 451

Which of the following technologies can successfully back up files that are used by other processes without stopping those processes?

A. Differential
B. Archive
C. Synthetic full
D. Open file

Answer: D

Explanation:

Comprehensive Detailed Explanation with all CompTIA Server+ SK0-005 Study Guide Reference:AnOpen file backupsolution allows files that are in use by other processes to be backed up without stopping those processes. This is especially important in environments where critical applications and services must remain operational while backup operations occur. Technologies like
VSS (Volume Shadow Copy Service) in Windows provide this functionality, enabling the backup of

open files such as databases or system files without disrupting the processes.

DifferentialandArchivebackups are types of backup strategies but do not specifically address backing up open or in-use files.

Synthetic fullbackups refer to a method of creating a full backup by combining incremental backups and a full backup, but it does not specifically address open files.

Question: 452

A datacenter has ten 40U racks in a hot/cold aisle configuration. The room has adequate air conditioning, but servers located near the top of the racks are shutting down due to issues with heat. Which of the following should be used to reduce issues with heat?

A. Rack balancing
B. Higher-capacity PDUs
C. Rail kits
D. Blanking panels

Answer: D

Explanation:

Comprehensive Detailed Explanation with all CompTIA Server+ SK0-005 Study Guide Reference:Blanking panelsare used to fill empty spaces in server racks. By blocking airflow gaps between servers, blanking panels help to direct airflow properly, ensuring that the cold air from the cold aisle is not wasted and that the hot air from the hot aisle is contained. This is crucial for maintaining the efficiency of cooling systems in a hot/cold aisle configuration.

Rack balancingaddresses the physical distribution of equipment, but it does not directly address thermal issues.

Higher-capacity PDUs(Power Distribution Units) provide more power, but they do not address cooling or airflow problems.

Rail kitsare used for mounting servers and are unrelated to cooling efficiency.

Question: 453

A Windows server has experienced a BSOD, and the administrator needs to monitor thereboot. The server is in a datacenter with no OOB management. Which of the following tools should the administrator use to complete this task?

A. Crash cart
B. IP KVM
C. Virtual administration console
D. Serial connection

Answer: A

Explanation:

Comprehensive Detailed Explanation with all CompTIA Server+ SK0-005 Study Guide Reference:ACrash cartis a physical console that can be connected directly to the server's video output and keyboard/mouse ports. This is useful in situations where remote management (OOB) is unavailable, allowing the administrator to directly view the server's display and interact with the system during the reboot process.

IP KVM(Keyboard, Video, Mouse) is used for remote management, but the scenario specifically mentions no OOB management.

Virtual administration consoleandSerial connectionscan be used for remote management or console access, but since OOB management is unavailable, a physical crash cart is the appropriate choice.

Question: 454

Which of the following should a server administrator configure to ensure access to motherboard software is restricted?

A. Biometric lock
B. Data-at-rest encryption
C. BIOS password
D. Bootloader authentication

Answer: C

Explanation:

Comprehensive Detailed Explanation with all CompTIA Server+ SK0-005 Study Guide Reference:ABIOS passwordrestricts access to the motherboard's firmware settings and can prevent unauthorized users from altering critical system configurations, such as boot sequence or hardware settings.

Biometric lockis a physical security feature but does not restrict access to motherboard software directly.

Data-at-rest encryptionprotects data stored on the system's hard drives but does not secure the BIOS or firmware.

Bootloader authenticationensures that the operating system's boot process is secured, but it is not a method to restrict access to the motherboard's software.

Question: 455

A Linux server was recently updated. The server now stops during the boot process with a blank screen and an fs> 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:

Comprehensive Detailed Explanation with all CompTIA Server+ SK0-005 Study Guide Reference:The fs> prompt typically indicates that the Unified Extensible Firmware Interface (UEFI) has entered the EFI shell due to a failure in locating the bootloader or a critical Linux boot file. This often occurs after an update where boot files are moved, renamed, or deleted. To resolve:

Verify the bootloader configuration (e.g., GRUB) and ensure the correct partition is set for booting. Recreate or repair the bootloader files if necessary.

Ensure the boot sequence in UEFI settings points to the correct disk and partition.

Other options:

A . USB flash drive booting:Unlikely unless the boot order was unintentionally altered.

C . BIOS not finding a bootable hard disk:Would generally result in a different error message.

D . BIOS firmware upgrade:Not typically required for this specific issue.

Question: 456

The internal IT department policy for monitoring server health requires that evidence of server reboots be collected and reviewed on a regular basis. Which of the following should be monitored to best provide this evidence?

A. Uptime
B. IOPS
C. CPU utilization
D. Event logs

Answer: D

Explanation:

Comprehensive Detailed Explanation with all CompTIA Server+ SK0-005 Study Guide Reference:Event logsare the primary source for tracking server reboots and other system events. Logs such as the system log in Linux (/var/log/syslog or /var/log/messages) or the Event Viewer in Windows provide timestamps and details about shutdowns, restarts, and boot processes. Regular monitoring ensures that unplanned reboots or system errors are identified and addressed.

Uptime:Indicates how long the server has been running since the last reboot but does not provide detailed reboot evidence.

IOPS:Refers to input/output operations per second, unrelated to reboots.

CPU utilization:Reflects processor activity but offers no reboot evidence.

Question: 457

When a user plugs a USB drive into a workstation and attempts to transfer a file from the server, an error appears that indicates the file cannot be copied. However, the user can copy or create new files from the local workstation. Which of the following is part of this security strategy?

A. HIDS
B. SIEM
C. ACL
D. DLP

Answer: D

Explanation:

Comprehensive Detailed Explanation with all CompTIA Server+ SK0-005 Study Guide Reference:Data Loss Prevention (DLP)policies are designed to prevent sensitive data from being copied, moved, or accessed in unauthorized ways. In this scenario, the DLP system is likely configured to block file transfers from the server to external storage devices, such as USB drives, while still allowing local file creation or movement.

HIDS (Host-Based Intrusion Detection System):Monitors system activities for malicious behavior but does not block file transfers.

SIEM (Security Information and Event Management):Provides centralized monitoring and analysis of security events but does not directly enforce file transfer restrictions.

ACL (Access Control List):Manages permissions but does not control data transfer policies.

Question: 458

An administrator is upgrading a group of legacy blade servers to the latest version of the hypervisor. After restarting the server, the upgrade is interrupted. The hypervisor is not available, and the update is not successful. Which of the following should the administrator do next to accomplish this upgrade?

A. Change the upgrade to a new installation and overwrite the boot drive.
B. Upgrade the blade server's firmware to support the new hypervisor.
C. Increase the blade server's onboard RAM.
D. Upgrade the blade server's hardware.

Answer: B

Explanation:

Comprehensive Detailed Explanation with all CompTIA Server+ SK0-005 Study Guide Reference:Firmware incompatibility is a common issue when upgrading hypervisors, especially on legacy hardware. Upgrading the firmware ensures that the hardware is compatible with the new hypervisor version. This includes updating BIOS, RAID controllers, and network adapters to the latest supported versions.

A . New installation:Risks data loss and is unnecessary if the issue is firmware-related.
C . Increasing RAM:May enhance performance but does not resolve compatibility issues.
D . Hardware upgrade:A last resort if firmware updates do not resolve the problem.

Question: 459

A server administrator set up a monitoring application on various servers to keep track of CPU usage, memory consumption, and disk space utilization. Which of the following should be configured on the application so it can send email alerts about a specific issue?

A. Event logs
B. Thresholds
C. Disk quotas
D. Uptime lengths

Answer: B

Explanation:

Comprehensive Detailed Explanation with all CompTIA Server+ SK0-005 Study Guide

Reference:Configuringthresholdsallows the monitoring application to define limits (e.g., CPU usage >90%, disk space < 10%) that trigger email alerts when exceeded. This proactive measure helps administrators address issues before they escalate.
Event logs:Provide historical data but do not send alerts.
Disk quotas:Restrict user disk usage but are unrelated to alerts.
Uptime lengths:Measure server runtime but do not trigger alerts.

Question: 460

A systems administrator notices the fans are running at full speed in a newly upgraded server. Which of the following should be done to address this issue?

A. Replace the PSU.
B. Reseat the video card.
C. Reseat the RAM.
D. Reattach the heat sink.

Answer: D

Explanation:

Comprehensive Detailed Explanation with all CompTIA Server+ SK0-005 Study Guide Reference:If the fans are running at full speed, it may indicate that the CPU or system temperature is too high. This is often caused by a poorly attached or improperly seated heat sink, which prevents efficient heat dissipation. Reattaching the heat sink with proper thermal paste resolves the issue.

A . Replace the PSU:Not relevant to temperature control.

B . Reseat the video card:Would not affect fan speed unless a GPU temperature issue exists, which is not stated here.

C . Reseat the RAM:Unrelated to fan operation.

Question: 461

Following a recent power outage, a server on a server farm has been running slowly. A technician receives a ticket with a request to check for anything out of the ordinary. The technician sees an orange light on the front console of the server, indicating an error. Which of the following should the technician check next?

A. CPU
B. Power supply
C. RAM
D. RAID battery

Answer: B

Explanation:

Question: 462

A server administrator is tasked with ensuring that information stored on company database systems is secure from bad actors. Which of the following security practices would accomplish this task?

A. Encryption at rest
B. Audit logging
C. BIOS password
D. Signal blocking

Answer: A

Explanation:

Question: 463

A senior administrator was notified that a junior administrator installed software on a server that should not have new software installed on it. The senior administrator decides to restrict the server via a Group Policy. Which of the following can be used to elevate the permissions?

A. Computer management
B. AD
C. Local users and groups
D. UAC

Answer: C

Explanation:

Question: 464

A technician is investigating a server's configuration to secure it from physical threats. The technician checks the BIOS and sees the following boot order:
USB Drive CDROM Drive Network Adapter RAID Controller
Additionally, the server's host OS is stored on mirrored drives. Which of the following should the technician do to secure the server?

A. Disable all except RAID Controller.
B. Block USB ports on the server.
C. Set BIOS password.
D. Move RAID Controller to the top of the list.

Answer: B

Explanation:

Question: 465

A systems administrator is setting up a second VLAN for end users. Which of the following should be provisioned for a DHCP server to be able to receive client requests from a different subnet?

A. DHCP options B. NIC teaming C. Zone transfers D. Relay agent E. VLAN ID

Answer: D

Explanation:

Question: 466

A certain application initially uses 1TB of drive space, but this is expected to double each year for the next two years. Which of the following is the minimum number of 1TB drives that are needed in a RAID 5 configuration?

A. 3

B. 4

C. 5

D. 6

Answer: C

Explanation:

RAID 5 is a storage configuration that uses striping with parity, providing both improved performance and fault tolerance. It requires a minimum of three disks, where data and parity information are distributed across all drives. The storage capacity of a RAID 5 array is calculated as (N - 1) * S, where N is the number of drives, and S is the size of each drive.

Storage Requirements:

Initial Storage: 1TB

After 1 Year: Doubles to 2TB

After 2 Years: Doubles again to 4TB

To accommodate 4TB of data in a RAID 5 setup, we use the formula:

$(N - 1) * 1TB \geq 4TB$

Solving for N:

$N - 1 \geq 4$

$N \geq 5$

Therefore, a minimum of 5 drives, each 1TB in size, is required to meet the projected storage needs. This configuration will provide a total usable capacity of 4TB, with 1TB allocated for parity to ensure fault tolerance.

Reference:

CompTIA Server+ Certification Exam Objectives (SK0-005): RAID Levels and Types

CompTIA Server+ (SK0-005) Study Guide: Chapter on Storage Solutions

Question: 467

Which of the following backup types only backs up files that have changed since the last full backup?

A. Differential
B. Open file
C. Incremental
D. Snapshot

Answer: A

Explanation:

Understanding different backup types is crucial for effective data protection strategies. Here's a breakdown of the relevant backup methods:
Full Backup: Captures all data, regardless of previous backups.
Differential Backup: Backs up data that has changed since the last full backup.
Incremental Backup: Backs up data that has changed since the last backup, whether it was full or incremental.
Snapshot: Captures the state of a system at a specific point in time.
A Differential Backup starts with a full backup. Subsequent differential backups save copies of all files that have been modified since that full backup. This means each differential backup includes all changes made since the last full backup, leading to larger backup sizes over time but faster restoration, as only the last full backup and the latest differential backup are needed.
In contrast, an Incremental Backup also begins with a full backup, but each subsequent backup only includes data that has changed since the most recent backup (whether full or incremental). This approach results in smaller backup sizes and quicker backup processes. However, restoration can be slower and more complex, as it requires the last full backup and all subsequent incremental backups to fully restore data.
Therefore, the correct answer is A. Differential, as it specifically refers to backing up files that have changed since the last full backup.
Reference:
CompTIA Server+ Certification Exam Objectives (SK0-005): Backup Methods
CompTIA Server+ (SK0-005) Study Guide: Chapter on Security and Disaster Recovery

Question: 468

A security team must ensure that unauthorized individuals are unable to tailgate through the data center's entrance. Which of the following should be implemented to stop this type of breach from happening?

A. Access control vestibule
B. Cameras
C. Bollards
D. Biometrics

Answer: A

Explanation:

Tailgating is a security breach where an unauthorized person follows an authorized individual into a secured area without proper credentials. To prevent tailgating, implementing an access control vestibule, also known as a mantrap, is highly effective.

An access control vestibule is a small room with two interlocking doors, where the first door must close and lock before the second door can be opened. This setup ensures that only one person can enter the secured area at a time, effectively preventing unauthorized individuals from gaining access by following someone else.

Other Options:

B . Cameras: While surveillance cameras can monitor and record entrances, they do not physically prevent tailgating. They serve as a deterrent and provide evidence after an incident has occurred but
do not stop unauthorized entry in real-time.

C . Bollards: Bollards are physical barriers used to prevent vehicular access to certain areas. They are not effective in preventing unauthorized individuals from tailgating into a building or secured area.

D . Biometrics: Biometric systems (e.g., fingerprint or retina scanners) verify the identity of individuals but do not address the issue of someone following an authorized person through an open door. Without physical barriers like an access control vestibule, tailgating can still occur even with biometric systems in place.

Therefore, the most effective measure to prevent tailgating is the implementation of an access control vestibule.

Reference:

CompTIA Server+ Certification Exam Objectives (SK0-005): Security and Disaster Recovery – Physical Security Concepts

Question: 469

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

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

Answer: B

Explanation:

Data Loss Prevention (DLP) systems are designed to monitor and protect sensitive data from unauthorized access, use, or transmission. By enabling DLP on a fileshare, the system administrator can track how data is accessed and used over time.

After allowing the DLP system to run for a week, the administrator can audit the permissions of the fileshare. This involves reviewing which users and groups have access to the fileshare and determining if their permissions align with their roles and responsibilities. Auditing permissions helps identify any discrepancies or excessive privileges that could lead to potential data breaches or unauthorized data exposure.

Other Options:

A . Manage the fileshare from an RDP session: Remote Desktop Protocol (RDP) allows administrators to remotely manage servers and their resources. While this is a method to access the server, it doesn't directly relate to auditing or managing fileshare permissions.

C . Audit the access to the physical fileshare: Auditing physical access involves reviewing who has physical entry to the hardware where the fileshare resides. While important, enabling DLP focuses on monitoring digital access and data movement rather than physical security.

D . Manage the permissions from the fileshare: Managing permissions involves setting or modifying user access rights.

However, before making changes, it's crucial to audit existing permissions to understand the current access control structure.

Therefore, after running DLP for a week, the appropriate action is to audit the permissions of the fileshare to ensure that access controls are properly configured and align with the principle of least privilege.

Reference:

CompTIA Server+ Certification Exam Objectives (SK0-005): Security and Disaster Recovery – Explain data security risks and mitigation strategies

Question: 470

A video card in an application server was recently replaced. The server administrator is now able to log in locally and view the screen with no issues; however, the administrator notices other performance issues, including the server's slow response time. The administrator reboots the server, but the issues persist. The server's cooling fans are running as normal, and the BIOS shows the dual power supplies are each working at 30%. Which of the following is most likely causing the performance issues?

A. A cable on the motherboard became unseated.
B. The firmware is incompatible with the video card.
C. A PSU is having power issues.
D. A CMOS battery failed.
E. Environmental temperature issues are being experienced.

Answer: B

Explanation:

Replacing hardware components, such as a video card, can sometimes lead to compatibility issues, especially if the system's firmware (BIOS/UEFI) does not support the new hardware. In this scenario, the server operates, and the display functions correctly, but there are performance issues like slow response times. This suggests that while the basic functionality is intact, there may be underlying compatibility problems affecting overall system performance.

Question: 471

Which of the following backup types is used to capture all data regardless of any changes from the previous backup jobs?

A. Incremental
B. Differential
C. Archive
D. Snapshot
E. Full

Answer: E

Explanation:

A Full Backup involves copying all data from a system, regardless of whether the data has changed

since the last backup. This process ensures that a complete, standalone copy of all data is available, facilitating straightforward restoration without the need for additional backup sets.

Other Backup Types:

A . Incremental Backup: This method backs up only the data that has changed since the last backup of any type (full or incremental). While it is storage-efficient and has a faster backup time, restoration requires all incremental backups since the last full backup, making the recovery process more timeconsuming.

B . Differential Backup: This approach backs up all data that has changed since the last full backup. It simplifies the restoration process compared to incremental backups, as only the last full backup and the latest differential backup are needed. However, it consumes more storage space and takes longer to perform as time progresses.

C . Archive: Archiving involves moving data that is no longer actively used to a separate storage system for long-term retention. It is not a backup type but a data management process aimed at freeing up primary storage space.

D . Snapshot: A snapshot captures the state of a system at a specific point in time. While it provides a quick way to restore a system to a previous state, it is typically stored on the same system and is not a substitute for a full backup.

In summary, the Full Backup is the only method that captures all data irrespective of changes, ensuring a comprehensive and independent data copy.

Reference:

CompTIA Server+ Certification Exam Objectives (SK0-005): Security and Disaster Recovery – Explain the importance of backups and restores

Question: 472

Which of the following benefits of virtualization is most important during a data center migration?

A. Portability
B. Orchestration
C. Overprovisioning
D. Scalability

Answer: A

Explanation:

During a data center migration, the ability to move workloads seamlessly from one environment to another is crucial. Portability in virtualization refers to the capability to transfer virtual machines (VMs) across different physical hosts or data centers with minimal disruption. This flexibility simplifies the migration process, reduces downtime, and ensures business continuity.

Other Options:

B . Orchestration: This involves the automated arrangement, coordination, and management of complex computer systems, middleware, and services. While orchestration enhances efficiency in managing virtual environments, it is not the primary benefit during a migration process.

C . Overprovisioning: This refers to allocating more resources than are physically available, relying on the assumption that not all resources will be used simultaneously. Overprovisioning can lead to resource contention and is generally avoided in well-managed virtual environments.

D . Scalability: This is the ability to increase or decrease resources and workloads dynamically based on demand. While scalability is a significant advantage of virtualization, it does not directly address the challenges associated with data center migration.

Therefore, Portability is the most critical benefit of virtualization during a data center migration, as it facilitates the efficient and reliable transfer of workloads between environments.

Reference:

CompTIA Server+ Certification Exam Objectives (SK0-005): Server Administration – Summarize the purpose and operation of virtualization

Question: 473

A corporation is implementing a cloud model that must meet security and privacy requirements for application development and testing. The company would like to use its own hardware to save money. Which of the following cloud models should be deployed?

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

Answer: D

Explanation:

A Private Cloud is a cloud computing model where the infrastructure is dedicated solely to a single organization. This model offers enhanced security and privacy, as all resources are behind the organization's firewall and exclusively managed by the organization. By utilizing their own hardware, the company can control costs and tailor the environment to meet specific security and compliance requirements, making it ideal for application development and testing.

Other Options:

A . Public: In a public cloud, services are delivered over the public internet and shared among multiple organizations. While cost-effective, it offers less control over security and privacy, which may not align with the company's requirements.

B . Community: This cloud model is shared among several organizations with similar interests and requirements. Although it provides more control than a public cloud, it still involves shared resources, which may not meet the company's strict security and privacy needs.

C . Hybrid: A hybrid cloud combines elements of both private and public clouds, allowing data and applications to be shared between them. While it offers flexibility, managing security and privacy across different infrastructures can be complex and may not provide the level of control the company desires.

Therefore, deploying a Private Cloud allows the company to utilize its own hardware, ensuring cost savings while meeting stringent security and privacy requirements for application development and testing.

Reference:

CompTIA Server+ Certification Exam Objectives (SK0-005): Server Administration – Summarize the purpose and operation of virtualization

Question: 474

An administrator is unable to get server updates from a WSUS server. However, the administrator was able to patch the server successfully last month. Which of the following should the administrator check first?

A. The software firewall on the local server
B. Stopped services on the local server

C. The primary DNS IP address assigned to the server
D. The hosts file on the local server

Answer: A

Explanation:

When a server that previously received updates from a Windows Server Update Services (WSUS) server suddenly fails to do so, one common cause is that the local software firewall is blocking the necessary ports or services required for WSUS communication. Firewalls can be reconfigured, either intentionally or accidentally, leading to such connectivity issues.

Reference:
CompTIA Server+ Certification Exam Objectives (SK0-005): Server Maintenance – Explain patching and update management

Question: 475

A systems administrator is selecting an authentication system for a data center. The company's security policy requires that the system support MF
A. Which of the following options should the administrator deploy to meet the policy requirements?
A. A retinal scan and a fingerprint reader
B. A key fob and an employee badge
C. An RFID chip and a PIN code
D. An alphanumeric, case-sensitive password with symbols

Answer: C

Explanation:

Multi-Factor Authentication (MFA) requires the use of two or more different authentication factors to verify a user's identity. The three main authentication factors are:
Something You Know – Password, PIN, security questions
Something You Have – Smart card, key fob, RFID chip
Something You Are – Biometrics such as fingerprint, retina scan
Option C (RFID chip and PIN code) meets the MFA requirement because:
RFID chip (Something You Have) provides a physical security token.
PIN code (Something You Know) adds a knowledge-based authentication factor.
Other Options:
A . Retinal scan and fingerprint reader: Both are biometric factors (Something You Are) and do not satisfy MFA requirements, which need at least two different categories.
B . Key fob and employee badge: Both are physical items (Something You Have), failing to meet MFA requirements.
D . An alphanumeric, case-sensitive password with symbols: This only represents Something You Know, not MFA.
Thus, RFID chip and PIN code (Option C) provides two separate authentication factors and meets the MFA policy

requirements.
Reference:
CompTIA Server+ Certification Exam Objectives (SK0-005): Security and Disaster Recovery – Explain authentication and access control

Question: 476

A systems administrator noticed a new network card is not being recognized by the operating system. Which of the following is most likely the cause?

A. The network cable is unplugged.
B. The driver is not loaded.
C. The IP settings are incorrect.
D. PXE boot is enabled.

Answer: B

Explanation:

When a newly installed network card is not being recognized by the operating system, the most likely cause is that the necessary driver is not installed or loaded. Device drivers act as intermediaries between the hardware and the operating system. Without the correct driver, the OS cannot communicate with or properly use the network card.
Other Options:
A . The network cable is unplugged: This would prevent network connectivity, but the OS would still recognize the network card itself.
C . The IP settings are incorrect: Incorrect IP settings can cause network issues but do not prevent the OS from recognizing the NIC.
D . PXE boot is enabled: PXE (Preboot Execution Environment) allows a system to boot from a network rather than a local disk. While PXE may affect booting, it does not directly cause a NIC to be unrecognized by the OS.
Thus, the most likely cause is that the driver for the network card is not installed or properly loaded. Reference:
CompTIA Server+ Certification Exam Objectives (SK0-005): Server Hardware – Explain server components and their functions

Question: 477

A technician is configuring a server with eight available drive bays. The technician wants a combination of the most redundancy and the maximum available storage. Which of the following RAID levels should the technician use?

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

Answer: D

Explanation:

When selecting a RAID level, both redundancy and storage efficiency must be considered:
RAID 1 mirrors data, providing high redundancy but reducing available storage to 50%. With eight drives, only four drives' worth of space would be available.
RAID 5 uses striping with distributed parity. It provides good storage efficiency (n-1) but can only tolerate one drive failure.
RAID 6 is similar to RAID 5 but uses dual parity, allowing two drives to fail while still maintaining data integrity. However, it reduces available storage to (n-2).
RAID 10 (also known as RAID 1+0) combines mirroring and striping. It provides excellent redundancy and performance, allowing multiple drive failures as long as they aren't in the same mirrored pair. Storage efficiency is 50%, but the redundancy and performance benefits outweigh this trade-off in high-availability environments.
Since the technician requires both maximum redundancy and high available storage, RAID 10 is the best choice.
Reference: CompTIA Server+ SK0-005 Official Textbook, Chapter 5 – Storage Solutions and RAID

Question: 478

A storage administrator is configuring a new array to provide storage to a virtual cluster. The array has a built-in backup agent to allow data that is stored on the array to be backed up. The storage administrator wants the virtual cluster to use a file-based storage protocol so the backup agent can just back up the files that change and not the whole data store. Which of the following storage protocols should the administrator use?

A. iSCSI
B. FC
C. FCoE
D. NFS

Answer: D

Explanation:
Different storage protocols operate at different levels (file-based vs. block-based):
iSCSI (Internet Small Computer Systems Interface) is a block-level protocol, meaning that the storage appears as raw disk space to the operating system. Backups would require copying entire blocks, rather than specific files.
FC (Fibre Channel) is another block-level protocol used in high-performance SANs. It does not support file-based storage natively.
FCoE (Fibre Channel over Ethernet) extends FC over Ethernet but still functions as a block-based storage solution.
NFS (Network File System) is a file-based storage protocol, making it the best option. It allows direct access to files, which enables backup software to back up only modified files instead of the entire block device.
Since the administrator requires a file-based storage protocol to ensure efficient incremental backups, NFS is the correct choice.
Reference: CompTIA Server+ SK0-005 Official Textbook, Chapter 6 – Storage Technologies

Question: 479

A server recently started sending error messages about running out of memory while in use. After a maintenance period during which more memory was added, the server is still unable to consistently remain powered on. Which of the following should the technician check first?

A. Memory compatibility
B. Memory speed
C. Memory slots
D. Boot order

Answer: A

Explanation:

When troubleshooting memory-related errors, it's essential to consider several key factors: Memory Compatibility: Memory modules must match the specifications supported by the motherboard (e.g., DDR4 vs. DDR5, ECC vs. non-ECC, buffered vs. unbuffered). Using incompatible RAM can cause boot failures, crashes, or instability.
Memory Speed: While RAM speed mismatches can affect performance, they typically do not cause a server to fail to power on. Most systems automatically adjust to the slowest module.
Memory Slots: Incorrect installation or damaged slots could be an issue, but the priority is to check compatibility first.
Boot Order: This setting controls the sequence in which devices boot the OS, which is unrelated to memory-related power issues.
Since the server remains unstable after adding memory, the most likely issue is an incompatibility with the motherboard. The first step is to verify whether the newly installed RAM meets the server's hardware requirements.
Reference: CompTIA Server+ SK0-005 Official Textbook, Chapter 4 – Server Components and Troubleshooting

Question: 480

A server in the data center reportedly has a BSOD and is not responding. The server needs to be rebooted. Which of the following is the best method to reboot the server remotely?

A. Crash cart
B. KVM
C. IRMI
D. shutdown -r

Answer: C

Explanation:

When a server crashes with a Blue Screen of Death (BSOD) and is not responding, the most effective way to reboot it remotely must be chosen:
Crash cart (A): A crash cart is a physical device used for troubleshooting at the server location. Since the administrator wants a remote reboot, this is not suitable.
KVM (B): A KVM (Keyboard, Video, Mouse) switch allows remote management of multiple servers but does not inherently provide reboot capabilities.

IRMI (C) - Intelligent Remote Management Interface: This is the best option. Many enterprise servers have Intelligent Platform Management Interface (IPMI) or Integrated Lights-Out (iLO) management tools that allow remote power control, including forced reboots.
shutdown -r (D): This command is used to restart a server via an operating system command-line interface. However, since the server is not responding, this command cannot be executed.
Since the server is unresponsive, the best method is to use the remote management interface (IRMI/IPMI/iLO) to force a reboot.
Reference: CompTIA Server+ SK0-005 Official Textbook, Chapter 7 – Remote Server Management

Question: 481

Which of the following would protect data in transit?

A. SSL/TLS
B. SHA-512
C. MD5
D. 3DES

Answer: A

Explanation:

Data in transit refers to data being transmitted over a network, and protecting it requires encryption: SSL/TLS (A): Secure Sockets Layer (SSL) and Transport Layer Security (TLS) are encryption protocols designed specifically to protect data in transit by encrypting communication between endpoints. This is the correct answer.
SHA-512 (B): Secure Hash Algorithm (SHA) is a hashing algorithm, not an encryption method. It is used for integrity verification rather than securing data in transit.
MD5 (C): Message Digest Algorithm 5 (MD5) is another hashing function, primarily used for checksums and integrity checks, not for encrypting transmitted data.
3DES (D): Triple Data Encryption Standard (3DES) is a symmetric encryption algorithm, but it is used for data at rest rather than data in transit.
Since the question specifically asks about protecting data in transit, SSL/TLS is the correct answer. Reference: CompTIA Server+ SK0-005 Official Textbook, Chapter 9 – Security and Encryption

Question: 482

An administrator is troubleshooting a performance issue with an application. The application performs many small reads and writes in its operation. The administrator determines that the underlying storage is currently configured to use a general-purpose array with RAID 5. Which of the following should the administrator consider doing in order to offer higher application performance yet still retain data protection?

A. Move the application to an array using RAID0.
B. Move the application to an array using RAID1.
C. Move the application to an array using RAID6.
D. Move the application to an array using RAID10.

Answer: D

Explanation:

RAID 5 is efficient for read-heavy workloads but is not well-suited for applications that perform many small reads and writes due to the overhead of parity calculations. To improve performance while maintaining redundancy:
RAID 0 (A): This would improve performance but does not offer any data protection, making it unsuitable.
RAID 1 (B): Provides redundancy but does not enhance write performance significantly.
RAID 6 (C): Offers better fault tolerance than RAID 5 but suffers from even higher write penalties due to double parity.
RAID 10 (D): Combines mirroring (RAID 1) and striping (RAID 0), significantly improving both read and write performance while ensuring redundancy. This makes it the best choice for applications with high-frequency small reads/writes.
Since the administrator needs higher performance and data protection, RAID 10 is the best choice. Reference: CompTIA Server+ SK0-005 Official Textbook, Chapter 5 – RAID Performance Considerations

Question: 483

A database developer recently requested a 4TB SATA drive utilizing spinning disks be added to an existing system for a new, heavily used database. The storage administrator turned down the request after reviewing it. Which of the following is the most likely reason the administrator did that?

A. The drive is not capable of deduplication.
B. The drive would have slow I/O performance.
C. The drive cannot use NTFS.
D. The drive is not compatible with the server OS.

Answer: B

Explanation:

A heavily used database requires high I/O (Input/Output) performance for frequent transactions. Considering the storage options:
The drive is not capable of deduplication (A): Deduplication is useful for saving storage space, but it is not a primary concern for database performance.
The drive would have slow I/O performance (B): Traditional spinning disk SATA drives have significantly lower IOPS (Input/Output Operations Per Second) compared to SSDs or enterprise-level NVMe storage. A high-traffic database benefits from low-latency, high-throughput storage, which SATA HDDs do not provide.
The drive cannot use NTFS (C): SATA drives can use NTFS, so this is incorrect.
The drive is not compatible with the server OS (D): Most modern operating systems support SATA drives.
Because SATA spinning disks cannot handle high transaction loads efficiently, the administrator likely rejected the request due to poor I/O performance.
Reference: CompTIA Server+ SK0-005 Official Textbook, Chapter 6 – Storage Technologies

Question: 484

A systems administrator is performing a routine update to a server. The administrator applies the update, restarts the server, and then conducts routine testing that reveals the critical functionality provided by the server is unavailable. Event logs indicate a core service is failing to start. The service is configured to start automatically, and rolling back the update

does not correct the issue. Which of the following is most likely causing the service failure?

A. The server requires another reboot to complete the rollback.
B. The administrator is not authorized to run the service.
C. The server requires further updates of other software components.
D. The account used to run the service has expired.

Answer: C

Explanation:

After an update, if a core service fails to start, potential causes include dependencies, authentication, or incomplete updates:

The server requires another reboot to complete the rollback (A): Some updates require multiple reboots, but if rolling back did not resolve the issue, a missing dependency is more likely.

The administrator is not authorized to run the service (B): If the administrator previously had access and the update broke the service, it is likely due to software conflicts rather than permissions.

The server requires further updates of other software components (C): Many services rely on dependent libraries, drivers, or patches. If an update replaces or removes a necessary component, the service may fail until additional updates are applied.

The account used to run the service has expired (D): Expired accounts usually cause authentication failures but not outright service crashes unless explicitly required for the update process.

Since rolling back the update did not fix the issue, the most likely reason is that additional updates or dependencies are missing.

Reference: CompTIA Server+ SK0-005 Official Textbook, Chapter 8 – System Maintenance and Troubleshooting

Question: 485

A technician is configuring a server that will need to accommodate a planned network upgrade. All hosts will be changed from 10Gb copper to 25Gb fiber. Which of the following would best suit the requirements?

A. FCoE
B. VLAN ID
C. SFP
D. HBA

Answer: C

Explanation:

The upgrade involves transitioning from 10Gb copper to 25Gb fiber, which requires appropriate network interface hardware:

FCoE (A) - Fibre Channel over Ethernet: This protocol allows Fibre Channel to run over Ethernet networks, but it does not address the physical change from copper to fiber.

VLAN ID (B): VLANs segment networks but do not impact the physical connectivity (copper vs. fiber). SFP (C) - Small Form-

factor Pluggable: SFP transceivers are hot-swappable modules used in network interfaces to support fiber connections. A 25Gb SFP28 transceiver would be required for the transition from 10Gb copper to 25Gb fiber. This is the correct answer.

HBA (D) - Host Bus Adapter: HBAs are typically used for storage connectivity (such as Fibre Channel SANs), not general network upgrades.

Since SFP modules allow servers to support fiber connections, this is the best choice for upgrading from 10Gb copper to 25Gb fiber.

Reference: CompTIA Server+ SK0-005 Official Textbook, Chapter 3 – Networking and Connectivity

Question: 486

Three new servers that use similar types of resources need to be set up, installed with the same OS, and ready within a short time frame. Which of the following installation types would most likely be used?

A. Physical clones
B. Bare metal
C. Virtualized
D. USB

Answer: C

Explanation:

Comprehensive and Detailed

When deploying multiple servers that require similar resources and need to be operational quickly, virtualization is often the most efficient approach. Virtualization allows for the creation of multiple virtual machines (VMs) on a single physical server, each running its own instance of an operating system. This method offers several advantages:

Rapid Deployment: VM templates can be used to quickly deploy standardized server environments, reducing setup time.

Resource Efficiency: Virtualization enables better utilization of hardware resources by allocating them dynamically among VMs.

Scalability: Additional virtual servers can be created as needed without the need for additional physical hardware.

Question: 487

A technician is installing a large number of servers in a data center with limited rack space. Which of the following would accomplish this goal while using the least amount of space?

A. Blade enclosure
B. Rack balancing
C. Rack mount
D. Rail kits

Answer: A

Explanation:

Comprehensive and Detailed

Blade enclosures are designed to maximize server density by housing multiple blade servers within a single chassis. This design minimizes the physical footprint compared to traditional rack-mounted servers. Key benefits include:
Space Efficiency: Multiple servers share power, cooling, and networking within a compact enclosure. Simplified Management: Centralized management of power, cooling, and networking simplifies administration.
Scalability: Easily add or replace blade servers without significant reconfiguration.

Question: 488

A technician receives reports that a file server is performing slower than usual after a power failure. While investigating the issue, the technician discovers the write cache was disabled. The technician checks a server configuration document and confirms the cache was previously enabled. Which of the following events most likely caused this change?

A. The cache hit rate was high
B. The volume reached 100% capacity
C. The RAID controller battery failed
D. A drive in the array was replaced

Answer: C

Explanation:
Comprehensive and Detailed
Write caching improves disk performance by temporarily storing write operations in faster cache memory before writing to disk. However, to prevent data loss during power failures, write caching relies on a battery-backed cache. If the RAID controller's battery fails, the system may automatically disable write caching to protect data integrity, leading to reduced performance.

Question: 489

Which of the following technologies best describes Hyper-V Replica?

A. High availability
B. VMotion
C. Live migration
D. Disaster recovery

Answer: D

Explanation:
Comprehensive and Detailed
Hyper-V Replica is a feature in Microsoft's Hyper-V virtualization platform that enables asynchronous replication of virtual machines to a secondary location. This technology is primarily designed for disaster recovery scenarios, allowing organizations to restore services in the event of site failures.

Question: 490

A server administrator needs to ensure proprietary data does not leave the company. Which of the following should the administrator implement to meet this requirement?

A. A review of users' activities

B. File-level monitoring
C. A SIEM solution
D. A DLP policy

Answer: D

Explanation:
Comprehensive and Detailed
Data Loss Prevention (DLP) policies are designed to detect and prevent unauthorized transmission of sensitive data outside the organization. Implementing DLP helps monitor and control data flows,

ensuring proprietary information remains within the company.

Question: 491

A recently deployed server is experiencing significant connection issues after being plugged into the network switch. An administrator was asked to review the server configuration and finds the following:
10Gb server NIC with 50-micron optic installed
10m single-mode fiber patch cable to the switch
10Gb switch interface with short-range optic installed
Which of the following steps should the administrator take next?

A. Replace the switch optic with one that supports MMF
B. Replace the server NIC with one compatible with single-mode fiber
C. Replace the optic in the server with a short-range optic
D. Replace the fiber cable with a 32.8ft (10m) OM4 cable

Answer: A

Explanation:
Comprehensive and Detailed
The server's NIC is equipped with a 50-micron optic, indicating it is designed for multimode fiber (MMF). However, the current setup uses a single-mode fiber (SMF) patch cable. To resolve the connectivity issue, the administrator should replace the switch's optic module with one that supports MMF, ensuring compatibility between the server NIC and the network infrastructure.

Question: 492

A technician notices a server that contains four fans is now louder than normal even though the temperature in the room has not changed and the load on the server has not increased. Which of the following is most likely the cause of the noise?

A. Defective fan
B. Change in UEFI
C. Unseated drive
D. Firmware upgrade needed

Answer: A

Explanation:
Comprehensive and Detailed
A sudden increase in fan noise without any changes in temperature or server workload is typically due to a defective fan.
If one fan fails, the remaining fans will automatically increase their speed to compensate for the loss, leading to increased noise levels.
Servers use intelligent fan control systems, which monitor temperature and fan health. If a fan stops functioning properly, the others will ramp up to maintain adequate cooling.
This issue can be resolved by inspecting and replacing the defective fan to restore normal cooling operation.
Reference:
CompTIA Server+ SK0-005 Official Study Guide
ExamTopics - CompTIA SK0-005

Question: 493

A technician is troubleshooting a database server that has 12 external disk arrays. Multiple disks are in a degraded state. The server is in a remote colocation facility and has on-site support. The technician notices all of the external storage arrays are degraded, but the internal arrays are functional. Which of the following steps should the technician take next to quickly address this issue?

A. Notify the users of the downtime and reboot the server from the RDP connection.
B. Notify the management team about the issue and travel to the colocation site with spare disks and RAID controllers.
C. Access the server with RDP and start the RAID utility software to rebuild the arrays.
D. Restore the data for the affected disks and notify the users of the estimated recovery time.

Answer: C

Explanation:
Comprehensive and Detailed
The external disk arrays being in a degraded state suggests a RAID failure or a communication issue with the storage system.
The quickest and most effective solution is to access the RAID utility software remotely via Remote Desktop Protocol (RDP) to diagnose and potentially rebuild the degraded arrays.
This avoids unnecessary downtime and delays caused by traveling to the site before attempting remote troubleshooting.
If the RAID utility shows disk failures, the technician can coordinate with on-site staff for further action, such as replacing failed drives or performing additional recovery steps.
Reference:
CompTIA Server+ SK0-005 Official Study Guide
ExamTopics - CompTIA SK0-005