DUMPSARENA

CompTIA Server+ Certification Exam

CompTIA SK0-005

Version Demo

Total Demo Questions: 54

Total Premium Questions: 549

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Server Hardware Installation and Management	110
Topic 2, Server Administration	135
Topic 3, Security and Disaster Recovery	182
Topic 4, Troubleshooting	122
Total	549

QUESTION NO: 1

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 the audit.log file size in the appropriate configuration file.

B. Decrease the duration of the log rotation cycle for the audit.log file.

C. Remove the log rotation directive from the audit.log file configuration.

D. Move the audit.log files to a remote syslog server.

ANSWER: A

Explanation:

Increase the audit.log file size in the appropriate configuration file. On Linux systems using the audit daemon, audit-log rotation is ordinarily controlled by `auditd`, using settings in `/etc/audit/auditd.conf`. In particular, `max_log_file` defines the maximum size, in megabytes, of an audit log file before the configured `max_log_file_action` is performed. When the action is set to `rotate`, increasing `max_log_file` allows the active `/var/log/audit/audit.log` file to accumulate more events before it is rotated. This directly reduces rotation frequency while continuing to retain audit records through the configured rotation and retention process. The related `num_logs` setting should also be sized appropriately for the organization's required retention period and available storage, so that rotated logs remain available rather than being discarded. Audit logging is security-sensitive, so the administrator should make this change through the service's supported configuration, validate the syntax, and reload or restart `auditd` under the organization's change-control procedures. The auditd configuration documentation describes `max_log_file`, rotation actions, and retention behavior: auditd.conf(5). General audit framework details are available in the Red Hat Enterprise Linux auditing documentation.

QUESTION NO: 2

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:

Account lockout threshold is the best policy to deter an online brute-force login attack because it limits the number of consecutive unsuccessful authentication attempts permitted for an account. Once that configured threshold is reached, the system locks the account for a defined period or until an administrator resets it. This directly disrupts automated password-guessing tools, which depend on being able to submit large numbers of credentials rapidly against the same account. By sharply restricting the available attempts, a lockout threshold makes exhaustive guessing impractical and provides a detectable security event that administrators can monitor and investigate.

In practice, the threshold should be balanced with a suitable lockout duration and monitoring controls to reduce the risk of unnecessary user lockouts while still slowing attackers. Organizations can supplement this policy with multifactor authentication, rate limiting, and alerting, but the account lockout threshold is the policy specifically designed to constrain repeated failed sign-in attempts. CompTIA's Server+ exam objectives include implementing account policies such as account lockout settings as part of identity and access management. Microsoft also documents that account lockout policies define the failed sign-in count that triggers a lockout and help mitigate password attacks. See Microsoft's Account Lockout Policy guidance and CompTIA exam objectives.

QUESTION NO: 3

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:

A single-mode fiber cable is used in place of multimode. is the most likely issue because the absence of link lights at both the server HBA and SAN switch points to a physical-layer connectivity problem. Fibre Channel links require compatible optical transceivers and cabling media at each end. Multimode Fibre Channel optics are designed for multimode fiber, whereas single-mode optics are designed for single-mode fiber. Installing single-mode fiber where a multimode link and optics are expected can create an unsupported or incompatible optical path, preventing the port from establishing link. Until the physical link initializes, the HBA cannot log in to the fabric, discover storage targets, or see presented LUNs. Fibre Channel deployments commonly use multimode optical fiber for shorter data-center runs, with the specific fiber type, wavelength, connector type, and transceiver all needing to match the supported link design. Cisco's transceiver guidance describes the importance of matching optical modules and supported cabling characteristics: Cisco Optical Transceiver Modules Compatibility. Additional Fibre Channel cabling and optical-interface information is available from the Fibre Channel Industry Association: FCIA Technology Overview.

QUESTION NO: 4

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 external firewall is blocking access and The default gateway is incorrect are the most likely causes because both directly affect the network path required to reach a package repository. A firewall positioned at the network perimeter can deny the repository's required outbound or inbound traffic, such as HTTPS traffic on TCP port 443. If that traffic is not permitted by policy, package-management clients cannot establish a connection to repositories located beyond that firewall. Firewall rules should therefore permit the necessary protocol, destination, and port while remaining as restrictive as possible.

The default gateway provides the route a server uses to reach destinations outside its local IP subnet. Package repositories are frequently hosted on another network or on the internet, so an absent or incorrect gateway prevents packets from reaching the appropriate router. Verifying the configured gateway, routing table, and end-to-end connectivity is a standard first step when repository access fails. Microsoft's networking guidance describes how the default gateway supplies the path to remote networks, while Red Hat documents the network connectivity prerequisites for accessing software repositories: Microsoft Learn: Network subnetting and default gateways and Red Hat: Managing software with DNF.

QUESTION NO: 5

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:

Close ports 69 and 1010 and rerun the scan. The system is intended to provide an internet-facing web application, so HTTP on port 80 and HTTPS on port 443 are the expected services that must remain reachable for users. In contrast, the scan identifies unexpected listening services on ports 69 and 1010. Their presence is especially significant alongside reports of poor performance and unsolicited pop-up content, which are indicators that an unauthorized service, unwanted software, or a compromise may be affecting the server. Port 69 is conventionally associated with Trivial File Transfer Protocol (TFTP), a protocol that has no built-in authentication or encryption and should not be exposed unless there is a specific, controlled operational requirement. The Internet Assigned Numbers Authority lists TFTP under port 69 in its service-name registry: IANA Service Name and Port Number Registry.

Closing both unexpected ports reduces the externally reachable attack surface and stops the suspicious services from accepting new connections. Rerunning Nmap afterward verifies that the remediation was applied and that only the required web services remain exposed. Nmap's documentation explains that a port reported as open has an application actively accepting connections, making such validation important: Nmap Port Scanning Basics.

QUESTION NO: 6

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 data. 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:

A drive failed in the OS array. The operating system resides on a two-drive RAID 0 array, which uses striping to improve performance and capacity but provides no fault tolerance or redundancy. Every file-system block is distributed across both member disks. Consequently, loss of either disk makes the complete OS volume unavailable, including essential boot files and system components. Once that drive failure occurred, the server could crash when it attempted to access missing portions of the operating system and would be unable to boot afterward because the RAID 0 logical volume could no longer be assembled into a usable volume. This precisely matches the observed sequence: a drive-failure alert followed by a crash and loss of boot capability. RAID configuration choices should reflect the criticality of the workload; boot volumes commonly use a redundant layout such as RAID 1 so that a single physical-disk failure does not make the server unavailable. Microsoft's overview of disk resiliency describes how striped volumes lack fault tolerance, while mirrored configurations provide redundancy. See Microsoft's Disk Management overview and CompTIA exam objectives resources.

QUESTION NO: 7

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:

A secure web server needs to accept encrypted web-client connections on port 443. This is the well-known service port registered for HTTPS, which protects HTTP traffic with TLS encryption. Allowing 443 enables users and applications to reach the server's secure website without exposing the web content through an unencrypted HTTP service. The Internet Assigned Numbers Authority lists HTTPS as using TCP port 443 in its service-name and port-number registry: IANA Service Name and Transport Protocol Port Number Registry.

The stated administrative requirement is Remote Desktop Protocol connectivity, so port 3389 must also be allowed for the administrator's authorized management connection. RDP uses TCP port 3389 by default; Microsoft documents this default listener port and explains that it can be changed only through explicit configuration. Therefore, permitting 3389 supports the

required remote graphical administration path, while permitting 443 supports the server's intended secure web-service role. See Microsoft's Remote Desktop listener port documentation. In a production deployment, both firewall rules should be scoped as tightly as practical—for example, restrict RDP source addresses to an administrative network or jump host—while still allowing the required HTTPS clients to access the web service.

QUESTION NO: 8

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:

Differential and **Incremental** backups are the appropriate selections because they avoid repeatedly copying all protected data, reducing backup-media consumption compared with performing a full backup for every backup cycle. A differential backup captures changes made since the most recent full backup. For recovery, the administrator needs the last full backup and only the latest differential backup, which keeps the restore process relatively fast while using less storage than repeated full backups. An incremental backup captures changes since the most recent backup, making it the most media-efficient method for frequent backup jobs. In a practical server backup plan, a full backup establishes the recovery baseline, while differential and incremental backup jobs protect subsequent changes without requiring a complete additional copy every time. This approach balances the organization's need to limit storage consumption with the need to maintain recoverable server data and reasonable restoration times. Microsoft describes the role of full and incremental backups in backup and recovery planning at Windows Server Backup documentation. Additional guidance on backup types and recovery considerations is available from the NIST contingency planning guide.

QUESTION NO: 9

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:

Social engineering is correct because the attacker is attempting to influence a person's behavior rather than initially defeating a technical control. Intentionally placing a USB drive where an employee is likely to find it relies on curiosity, helpfulness, or a desire to identify the owner. If the employee connects the device to a company system, the malicious code

can execute, install malware, or otherwise create a path into the environment. This specific technique is commonly called USB baiting or a USB-drop attack, and it is a physical form of social engineering.

The core security lesson is that removable media found in uncontrolled locations must be treated as untrusted. Organizations reduce exposure through user-awareness training, policies governing removable media, endpoint controls that restrict USB devices, and procedures for turning in unknown devices without connecting them to a system. CISA describes social engineering as tactics that exploit human interaction to gain access or information and advises users to be cautious with unexpected items and requests. See CISA: Avoiding Social Engineering and Phishing Attacks and CISA: Malware.

QUESTION NO: 10

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:

Overprovisioning is the practice of assigning virtual machines more virtualized resources—such as vCPUs, memory, or storage capacity—than the hypervisor's host has physically available at one time. This approach is possible because virtual workloads commonly do not all reach their maximum resource demand simultaneously. For example, an administrator may allocate a combined total of 64 vCPUs to several virtual machines running on a host with fewer physical processor cores, relying on the hypervisor to schedule access to the underlying CPU resources. Similarly, thin-provisioned virtual disks can present more logical storage to guests than has been physically committed in the storage pool. Effective overprovisioning requires monitoring utilization and performance so that resource contention does not create unacceptable latency or availability issues when multiple workloads demand peak capacity concurrently. VMware describes CPU overcommitment and its scheduling considerations in its vSphere Resource Management documentation. Microsoft also explains virtual-machine resource allocation and host capacity considerations in its Hyper-V scalability planning guidance.

QUESTION NO: 11

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 correct because it creates a controlled transition point between a general-access area and a restricted data-center space. Often called a mantrap, this vestibule uses two interlocking doors: the entry door must close and the person must be authenticated before the interior door can open. This arrangement helps ensure that only one authorized individual enters at a time and substantially reduces tailgating or piggybacking, where an unauthorized person follows an authorized user through a secure entrance. Authentication may be implemented with a badge reader, PIN, biometric factor, or a combination of factors, and the event can be logged for auditing and incident investigation. In a data

center, an access control vestibule is particularly appropriate at entrances to server rooms, equipment cages, or other zones containing critical infrastructure because it actively enforces access restrictions at the physical boundary. CompTIA includes physical security methods within the Server+ security domain, where candidates must apply controls suitable for protecting servers and their environments. NIST also identifies physical access control as a safeguard for limiting access to organizational facilities and systems. See the CompTIA exam objectives resource page and NIST SP 800-53 Rev. 5.

QUESTION NO: 12

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. RFI

ANSWER: A

Explanation:

A mantrap is a physical access-control vestibule with two interlocking doors: after a person enters through the first door, that door closes and locks before the second door can be opened. This design creates a controlled holding area in which the facility can validate that only an authorized individual is proceeding into a protected space. In this scenario, showing a driver’s license establishes identity at entry, while the scale adds an anti-tailgating or anti-piggybacking control. Comparing the measured weight to an expected range can help detect whether more than one person entered the enclosure or whether an unusual item is being carried through the checkpoint. Access to the second locked door is granted only after this validation step completes. Mantraps are especially appropriate for high-security data-center areas because they enforce single-person passage and provide a physical barrier between the public or lower-security area and the protected environment. This aligns with NIST guidance requiring physical access controls for facilities and controlled entry points, including mechanisms that manage and monitor access to sensitive areas. See NIST SP 800-53 Rev. 5 and CISA Physical Security.

QUESTION NO: 13 - (HOTSPOT)

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.

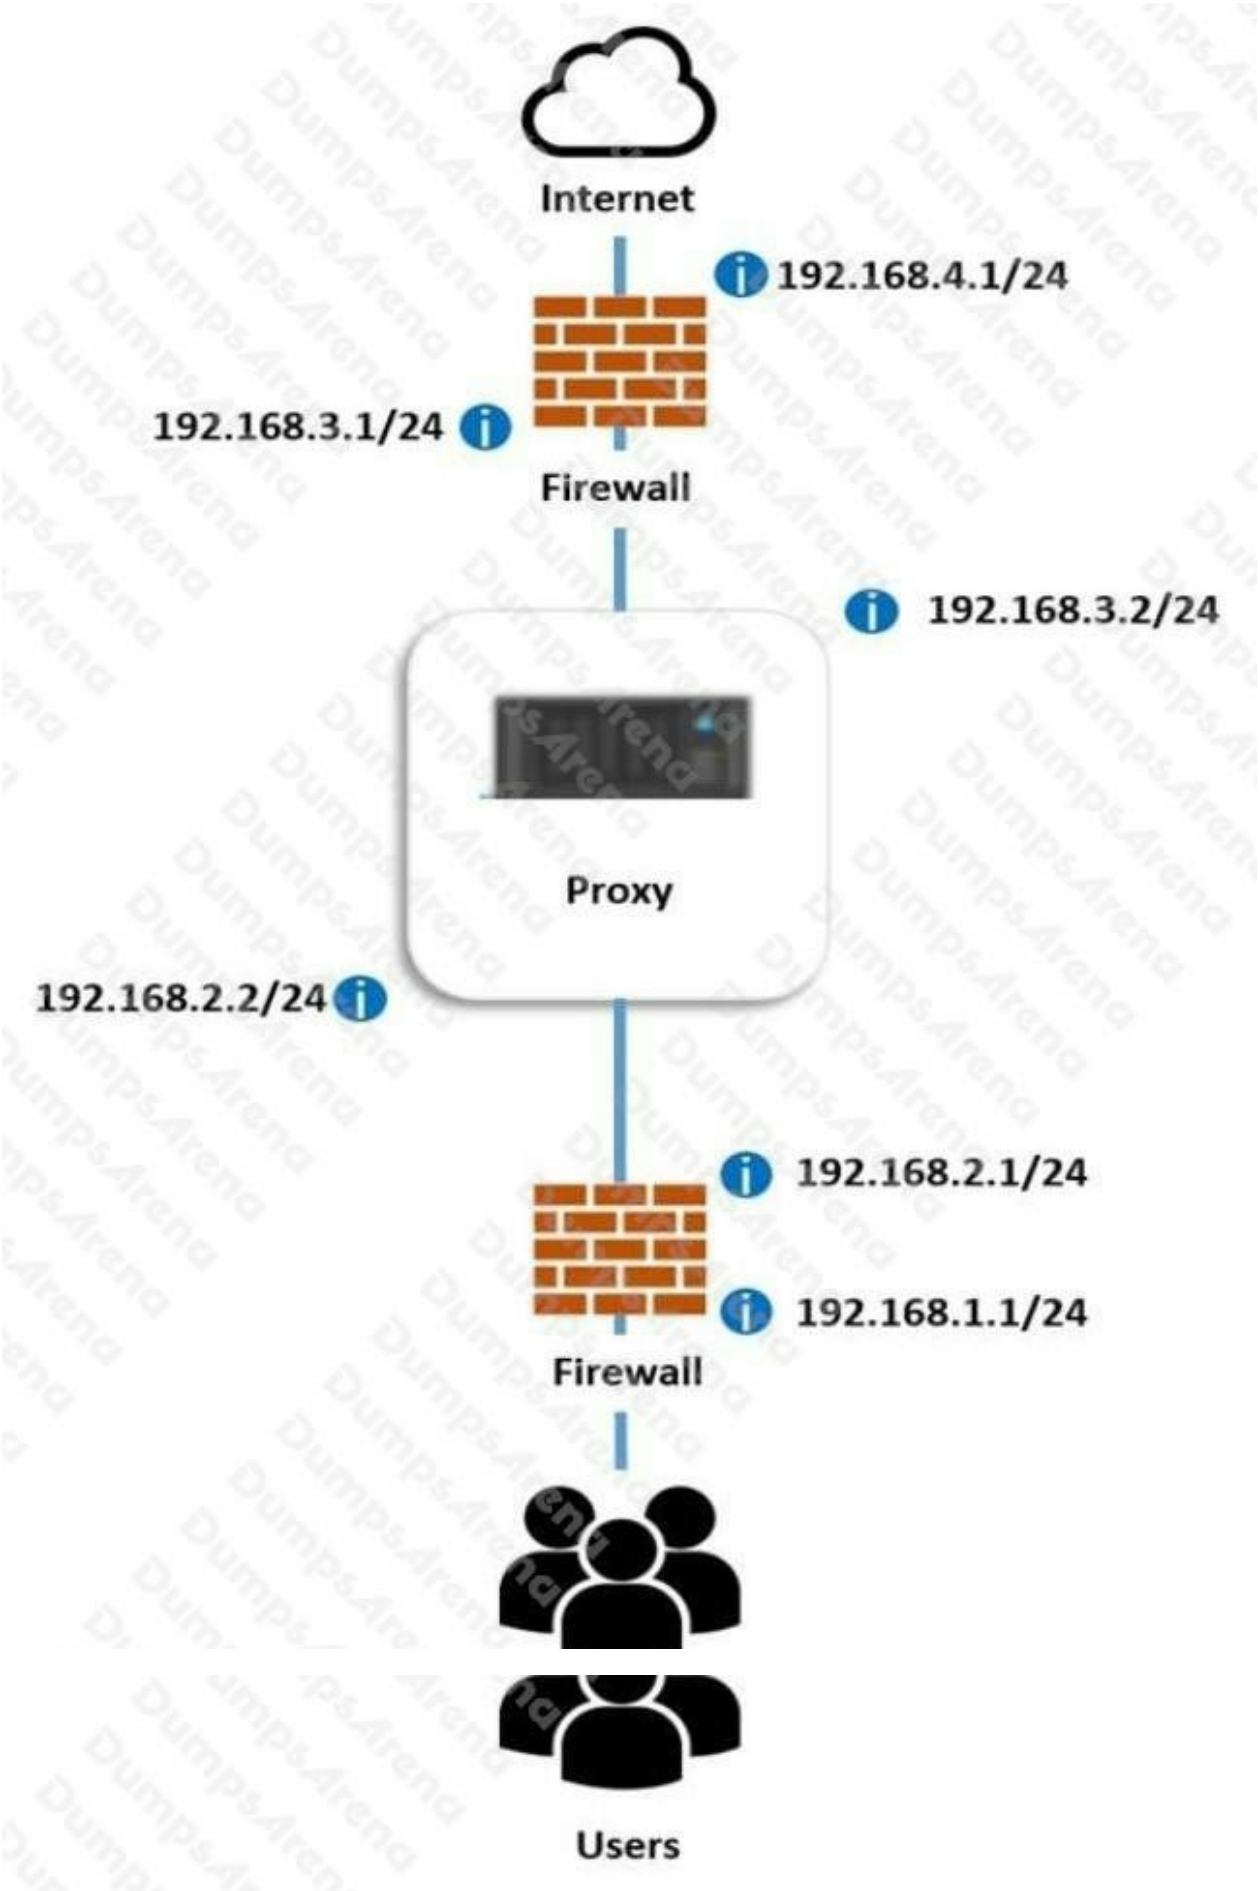
INSTRUCTIONS

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.

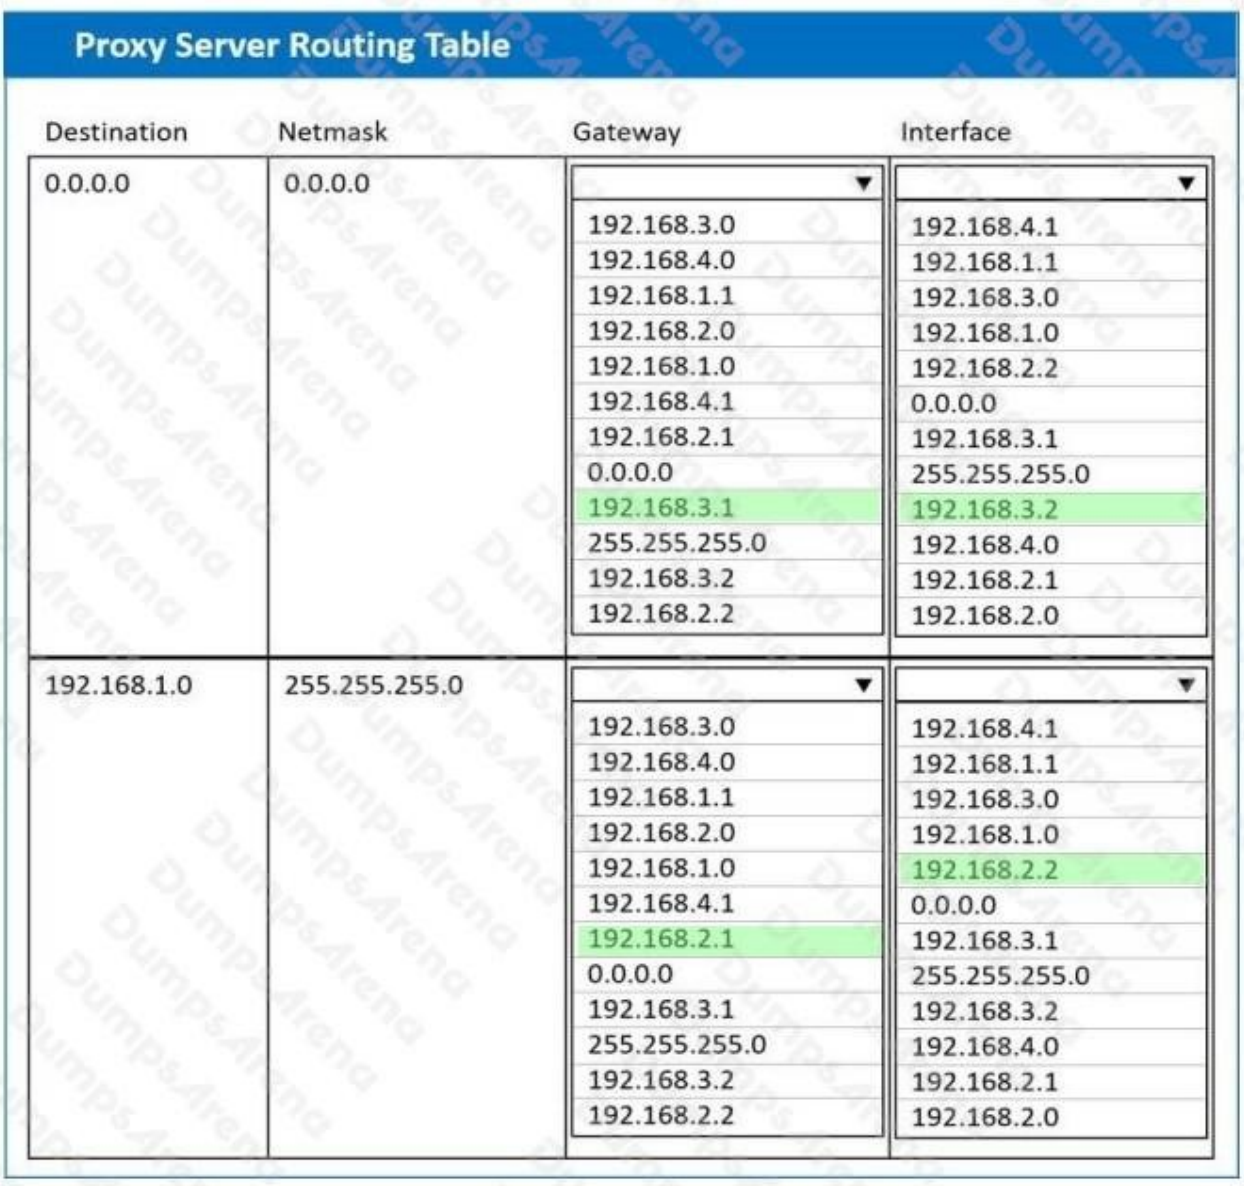

Internet
192.168.4.1/24
192.168.3.1/24
Firewall
192.168.3.2/24
Proxy
192.168.2.2/24
192.168.2.1/24
192.168.1.1/24
Firewall
Users

Hot Area:

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:

The correct routing configuration gives the proxy server a route in each required direction. The route for **192.168.1.0** with netmask **255.255.255.0** identifies the internal corporate subnet. Its next hop is **192.168.2.1**, and it leaves through the proxy interface at **192.168.2.2**. These addresses are on the same transit subnet, so the proxy can deliver response traffic back to the internal firewall, which can then forward it to corporate users on the 192.168.1.0/24 network. This explicit route is important because the proxy needs a known return path for sessions initiated by internal clients.

The route for destination **0.0.0.0** with netmask **0.0.0.0** is the default route. It forwards traffic for all networks not represented by a more-specific entry to gateway **192.168.3.1** through the proxy interface at **192.168.3.2**. Since that gateway is reachable on the internet-facing subnet, it is the proper next hop for Internet-bound proxy traffic. Routing selects the most specific matching route first, so traffic for 192.168.1.0/24 uses the internal route, while external destinations use the default route. Together, these entries provide the bidirectional network reachability required for users to access Internet resources through the new proxy server. Cisco's routing overview explains that a default route is used when no more-specific route matches: Cisco: Default Routing. Microsoft also documents the destination, netmask, gateway, and interface concepts used in host routing tables: Microsoft route command reference.

QUESTION NO: 14

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:

Remote desktop and **IP KVM** provide remote interactive administration of servers located across a data center. Remote Desktop Services enables an administrator to establish a graphical remote session to a Windows Server over the network. From that session, the administrator can perform routine maintenance tasks such as reviewing logs, changing configuration, applying updates, managing services, and troubleshooting operating-system-level issues without traveling to each rack. Microsoft documents Remote Desktop as a way to connect to and control a remote PC or server from another device.

An IP KVM provides network-based keyboard, video, and mouse access to attached physical servers. It gives an administrator console-like visibility and control remotely, which is especially valuable when a server has no working network configuration, cannot boot into Windows, requires BIOS or firmware changes, or needs observation during startup. Many IP KVM solutions also support remote virtual media, allowing authorized installation or recovery media to be mounted across the network. Together, these tools support both in-band Windows administration and out-of-band physical-console maintenance for all servers from a central location.

References: Microsoft Remote Desktop clients; Raritan KVM over IP.

QUESTION NO: 15

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:

216.176.128.10 best fits the scenario because it is formatted as a valid IPv4 unicast address and does not fall within an IPv4 range reserved for private internal addressing. A company that has been assigned a public address block can use those globally unique addresses on systems within its server farm. This is commonly appropriate when hosts or services need directly routable, Internet-reachable addressing, subject to the organization's routing, firewall, and security policies. Public address space is allocated through the Internet number-resource system, with registration and assignment coordinated by Regional Internet Registries. The Internet Assigned Numbers Authority maintains the special-purpose IPv4 registry, which identifies address blocks reserved for particular non-public uses; 216.176.128.10 is not in one of the private-use blocks defined by RFC 1918. RFC 1918 specifically reserves 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16 for private internets, leaving an address such as 216.176.128.10 suitable as a public-address example. See RFC 1918 and the IANA IPv4 Special-Purpose Address Registry.

QUESTION NO: 16

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 the components most commonly designed for hot-swapping in enterprise servers. A server with redundant, hot-swappable power supplies can continue receiving power from its remaining supply while a failed unit is removed and replaced. This capability minimizes service interruption, provided the server chassis, power configuration, and replacement unit support hot-swap operation. Dell describes redundant power supplies as enabling continued operation when one supply fails: Dell PowerEdge redundant power supply configuration.

The hard drive is also frequently installed in a hot-swap drive bay, especially in servers using RAID. The RAID controller and operating system can keep storage available through redundancy while an administrator replaces a failed drive. After insertion, the replacement drive can be recognized and used to rebuild the affected RAID array without shutting down the host. This depends on using a hot-swap backplane, compatible drive carrier, and a RAID level that provides fault tolerance. Microsoft documents the role of redundant storage configurations and rebuilding after a disk failure in Storage Spaces fault tolerance.

QUESTION NO: 17

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 appropriate because it reduces electromagnetic emissions that could be intercepted outside the facility. This control is commonly implemented through shielding techniques, such as a properly designed Faraday enclosure, shielded rooms, conductive building materials, and protected cable pathways. By attenuating radio-frequency and other electromagnetic signals at the facility boundary, signal blocking helps limit the ability of an external party to detect, monitor, or collect information from those emissions. The National Institute of Standards and Technology describes TEMPEST-related protections as controls intended to address compromising electromagnetic emanations; see NIST's TEMPEST glossary entry.

Camouflage complements signal blocking by making the data center and its potentially revealing external features less conspicuous. Concealing or blending elements such as communications infrastructure, equipment areas, and physical building characteristics reduces an observer's ability to identify the facility as a valuable transmission source or conduct effective reconnaissance against it. This is a form of deterrence and concealment within a defense-in-depth physical-security strategy. NIST notes that physical and environmental protections are an important part of safeguarding information systems and facilities; see NIST SP 800-53, Physical and Environmental Protection controls.

QUESTION NO: 18

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.

ANSWER: A C F

Explanation:

"Run a vulnerability scanner on the server," "Patch the vulnerability," and "Update the antivirus software" provide a layered remediation approach that directly addresses both the exploited weakness and the risk of a similar future compromise. A vulnerability scan identifies the database-software flaw and can reveal additional missing patches, insecure configurations, or exposed services requiring remediation. Patching the vulnerability is essential because it removes the known condition that enabled the infection; patches should be tested through the organization's change-management process and deployed promptly to reduce exposure while protecting service availability. Updating the antivirus software supplies current detection signatures and protection capabilities, helping the server detect, block, or quarantine malware associated with the incident and newer variants. Together, these actions support a practical security lifecycle: identify exposure, remediate the root cause, and strengthen endpoint detection and prevention. This approach is consistent with established guidance to remediate known vulnerabilities and keep security software current. See the CISA Known Exploited Vulnerabilities Catalog and Microsoft Defender Antivirus update guidance.

QUESTION NO: 19

Following a recent power outage, a server in the data center has been constantly going offline and losing its configuration. Users have been experiencing access issues while using

the application on the server. The server technician notices the date and time are incorrect when the server is online. All other servers are working. Which of the following would most

likely cause this issue? (Select two).

A. The server has a faulty power supply.

B. The server has a CMOS battery failure.

C. The server requires OS updates.

D. The server has a malfunctioning LED panel.

E. The servers have NTP configured.

F. CPU frequency scaling is set too high.

ANSWER: A B

Explanation:

“The server has a CMOS battery failure” is supported by the loss of configuration together with an incorrect date and time after the server loses power. The CMOS/RTC battery maintains the real-time clock and firmware configuration when external power is absent. When it is depleted or failed, an outage can cause the system clock to reset and firmware settings to be lost, potentially affecting services that depend on accurate time, such as authentication, certificates, scheduled tasks, and logging.

“The server has a faulty power supply” also fits the recurring offline condition. A failing power supply can produce intermittent power loss, unexpected shutdowns, or an inability to remain powered on after the initial outage. Because every other server remains operational, the issue is localized to this server rather than being a continuing facility-wide power or network condition. Together, a failed CMOS battery accounts for the reset clock and lost configuration, while a faulty power supply accounts for the server repeatedly going offline and the resulting application-access interruptions. Dell’s guidance on RTC battery replacement describes loss of date/time and BIOS settings as typical symptoms, and Microsoft documents that unexpected shutdowns can interrupt services and system operation. See Dell RTC battery guidance and Microsoft shutdown documentation.

QUESTION NO: 20

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 the best solution because it uses distributed dual parity across all four disks. With four disks, it provides usable capacity equivalent to two disks while retaining the ability to recover from the failure of any two disks in the array. This is the greatest fault tolerance provided by the listed standard RAID configurations when all four disks are used to create one data volume. Dual parity means that the RAID controller or software has sufficient parity information to rebuild data after a first disk failure and still protect the volume if a second disk fails before the first replacement and rebuild process are complete. That additional failure tolerance is particularly valuable for server data volumes, where larger-capacity disks can make rebuilds take longer and increase the importance of maintaining protection during that period. RAID 6 requires at least four drives, so the available disk count meets its minimum configuration requirement exactly. CompTIA’s Server+ objectives include RAID configuration and fault-tolerance concepts as core server-storage knowledge; see the CompTIA Server+ certification page. For a technical description of RAID 6 dual-parity protection and its ability to tolerate two drive failures, see the NetApp RAID overview.

QUESTION NO: 21 - (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



ANSWER:



Explanation:

The correct completed configuration moves VM Host 1 and Mail Relay out of Rack 1 and into the available Rack 2 positions, leaving VM Host 2 in Rack 1. This distributes server workload across the two racks instead of concentrating the equipment in the rack whose PDU A is already reading 13A. The move must be an actual relocation: each moved device appears only once, in Rack 2, rather than remaining displayed in Rack 1 as well.

After the relocation, use each device's pencil control to set its power-feed selections so that every device requiring redundant power has independent feeds from the rack's PDU A and PDU B. Redundant power means that the loss of one PDU, UPS output, or power path does not remove power from the device; the remaining feed continues supplying it. The PDU selections also have to balance the calculated maximum draw across the two PDUs in each rack. The final displayed maximum for every PDU must be no greater than 9.6A, which is 80 percent of the 12A PDU rating. This operating margin is important because startup current, workload changes, and UPS transfer conditions can otherwise push a fully loaded circuit into an overload condition.

This approach resolves both requirements at once: workloads are distributed between racks, and each rack has resilient, balanced A/B power paths that stay within the required design limit. CompTIA's Server+ exam objectives include power, cooling, and high-availability considerations. For practical redundant-power guidance, see Eaton's UPS basics overview.

QUESTION NO: 22

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:

Look for an unauthorized HTTP service on this server is the appropriate next action because successful connections to localhost on TCP ports 80 and 443 establish that processes on the server are listening for web traffic. Port 80 is the well-known port for HTTP, while port 443 is the well-known port for HTTPS, as listed by the Internet Assigned Numbers Authority (IANA): IANA Service Name and Port Number Registry. The technician should identify the listening process and validate that the web service is approved, expected, properly configured, and owned by the intended account. On Linux, commands such as `ss -ltnp` can associate listening TCP sockets with their processes, and the service manager can then be used to inspect the relevant service configuration and status. This investigation is especially important when a server is not intended to host web content, since an unexpected listener could expose an administrative interface, a default application, or unapproved software. Red Hat documents the use of `ss` for examining listening sockets and associated processes in its networking guidance: Using the ss utility. Confirming the HTTP/HTTPS service is therefore the direct, evidence-based follow-up.

QUESTION NO: 23

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:

The firmware is incompatible with the video card. The performance symptoms began immediately after the video-card replacement, making the newly installed hardware and its interaction with the server platform the most relevant troubleshooting area. A graphics adapter can provide basic local video output while still having compatibility problems with the server's installed BIOS/UEFI firmware, option ROM handling, PCIe initialization, or device resource allocation. This can leave the operating system usable but contribute to degraded responsiveness or instability after startup. A reboot does not resolve a persistent firmware-to-device compatibility condition because the same initialization process occurs again using the unchanged firmware and adapter combination.

The normal fan operation and the reported balanced 30% loading on both power supplies provide useful evidence that the server is not currently indicating an obvious cooling or power-capacity problem. The appropriate corrective path is to verify the server vendor's supported-adapter list, review release notes for BIOS/UEFI and video-card firmware, and apply an approved firmware update during a controlled maintenance window. Server firmware updates commonly include improved compatibility and initialization support for add-in hardware. See Dell BIOS update guidance and HPE firmware update best practices.

QUESTION NO: 24

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:

Core is the best installation method because it deploys a minimal server operating system footprint, installing only the components needed for core server roles and omitting the full local graphical management environment. Reducing installed features and background services lowers consumption of disk space, memory, and processor resources, leaving more capacity available for hosted applications and workloads. A smaller installation also has fewer components to maintain and patch, helping administrators reduce the operational and security exposure of the server while still supporting remote administration. In Windows Server, Server Core is specifically designed for this purpose: it provides a lightweight installation option that can run common infrastructure roles and be administered with PowerShell, Windows Admin Center, Remote Server Administration Tools, and other remote-management methods. The selection directly addresses the requirement to optimize server resources rather than simply choosing where the operating system is installed or a management-interface preference. Microsoft describes Server Core as a minimal installation option that includes most server roles while omitting the traditional desktop experience, reducing the servicing requirements and resource footprint. See Microsoft Learn: What is Server Core? and Microsoft Learn: Install Server Core.

QUESTION NO: 25

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:

Template is correct because it provides a standardized, preconfigured starting point for deploying virtual machines. An administrator can prepare a template with the common operating system configuration, virtual hardware settings, security baselines, updates, drivers, and required foundational software that the similar VMs need. New VMs can then be deployed from that template rather than built individually through repeated operating-system installation and configuration steps. This substantially reduces provisioning time while promoting consistency with the established system requirements.

A template also supports the stated need for VMs that have similar requirements but may run different applications. The shared baseline can be deployed first, and application-specific settings or software can be applied afterward. This approach helps ensure that each new VM begins from an approved, repeatable configuration and reduces the chance of configuration drift caused by manual, one-by-one builds. CompTIA Server+ includes virtualization deployment and administration concepts, where using standardized images/templates is a core operational practice. VMware describes VM templates as master copies that can be used to create and provision virtual machines: VMware vSphere Documentation. Microsoft similarly documents the use of generalized images to create multiple consistent virtual machines: Azure Compute Gallery documentation.

QUESTION NO: 26

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 correct because Windows Server can use dynamic disks to create software-based RAID volumes managed by the operating system rather than by a dedicated hardware RAID controller. After disks are converted to dynamic disks, Disk Management can use space across multiple physical disks to configure volume types that provide performance or fault-tolerance features, including striped, mirrored, and RAID-5 volumes where supported by the Windows Server edition. A mirrored volume stores duplicate copies of data on two disks, while a RAID-5 volume distributes data and parity information among three or more disks. This lets an administrator implement RAID functionality when appropriate hardware support is unavailable or when an OS-managed storage design is required. Dynamic disks are a legacy Windows disk-management technology; for newer deployments, Microsoft generally recommends Storage Spaces for many software-defined storage scenarios. Nevertheless, for the technology specifically associated with creating traditional Windows software RAID volumes in Disk Management, dynamic disks are the appropriate answer. See Microsoft's Disk Management overview and Microsoft's dynamic disk documentation.

QUESTION NO: 27

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 because the symptoms began after hardware was changed and are random, which strongly points to an intermittent hardware-level failure. RAM is fundamental to every operating-system and application workload. If a DIMM has a defect, is incompatible at the configured speed or timing, or returns corrupted data, Windows can encounter an unrecoverable kernel error and stop with a bug check (BSOD). Microsoft describes a bug check as occurring when Windows cannot continue operating safely after a serious error; the displayed stop code and crash dump can help confirm the failure pattern.

A DIMM that is not fully inserted or whose retaining clips are not secured can also produce intermittent electrical contact. This may allow the server to boot and run at times while generating unpredictable memory-access errors under load, vibration, or temperature changes. A practical remediation is to power down using approved maintenance procedures, reseat the upgraded DIMMs, verify that they are supported by the server platform, and run the platform's memory diagnostics. If errors persist, test modules individually or replace the failing module. See Microsoft's bug-check code reference and Windows kernel debugging documentation for crash-analysis context.

QUESTION NO: 28

A server administrator is preparing a maintenance window to replace a failed RAID controller in a production server. Which of the following should the administrator complete before beginning the replacement? (Choose two.)

A. Verify current backups

B. Document the existing RAID configuration

C. Initialize all member drives

D. Create a new logical volume

E. Format the existing volume

ANSWER: A B

Explanation:

Verify current backups and **document the existing RAID configuration** should be completed before replacing the controller. A verified backup provides a recovery option if the replacement process fails, if the storage configuration is damaged, or if an unexpected issue affects the logical volumes. Administrators should confirm that the backup completed successfully and that the protected data can be restored.

Documenting the existing RAID configuration preserves information such as the RAID level, member-drive order, logical-volume settings, controller firmware version, and cache configuration. This information is important because a replacement controller may need to import or reconstruct the existing array metadata correctly. Initializing the drives or creating a new array would risk overwriting existing configuration or data. NIST recommends maintaining backups and recovery procedures as part of contingency planning; see NIST SP 800-34 Rev. 1.

QUESTION NO: 29

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.

H. Configure RAID 10 for the application data.

ANSWER: B E H

Explanation:

Installing the OS on a RAID 1 array provides a dedicated, mirrored OS volume. Each member contains an identical copy of the operating system, so the OS remains available after a single member-disk failure. This also maintains separation between the OS storage and the application-data storage, limiting the impact of a data-volume issue.

Configuring RAID 10 for the application data is the appropriate design for an IOPS-sensitive production workload that also requires strong availability. RAID 10 combines mirroring with striping: striping distributes I/O across disks for high read and write performance, while mirrored pairs provide redundancy. It can continue operating through more than one disk failure when the failed disks are in different mirror pairs. This architecture is generally favored for transactional and other write-intensive workloads because it avoids the parity-calculation write penalty associated with parity RAID. Lenovo describes RAID 10's striping, mirroring, performance, and fault-tolerance characteristics in its RAID levels documentation.

Using SSD hard drives for the data application array further supports the IOPS requirement. SSDs have substantially lower latency and are capable of far more random I/O operations than spinning disks, making them well suited to application data such as databases and virtualized workloads. Microsoft also discusses SSD performance considerations in its disk performance documentation.

QUESTION NO: 30

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 correct because it centralizes log records from servers into a separate logging platform or repository. Once events are forwarded and retained centrally, administrators and auditors can search and review the records even when an

individual server is offline, failed, being rebuilt, or otherwise unavailable. This supports both operational troubleshooting and security investigations by preserving a broader, time-correlated view of activity rather than relying on locally stored logs on each server.

For this policy to be meaningful, log aggregation should include reliable forwarding from every server, timestamp synchronization, secure transport, and retention settings that meet the organization's review and compliance requirements. A central collector, SIEM, or log-management server can normalize and index received events, making records available independently of the originating system's current state. The National Institute of Standards and Technology describes centralized log management as a key capability for organizations that need to collect, store, and analyze logs from many systems. See NIST SP 800-92: Guide to Computer Security Log Management and Microsoft Azure Monitor Logs ingestion overview.

QUESTION NO: 31

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:

Bootloader password and **BIOS/UEFI password** provide complementary controls over the pre-operating-system startup path, which is the point at which an unauthorized user could otherwise boot installation media and replace or add an operating system. A bootloader password, such as a GRUB password on Linux systems, restricts access to bootloader editing and interactive commands. When correctly configured, it prevents an individual at the server console from altering boot parameters or selecting bootloader functions that could facilitate booting an unauthorized installer. GNU documents GRUB authentication and its use in protecting restricted boot menu entries and command-line access: GNU GRUB Manual: Authentication and authorisation.

A BIOS/UEFI password protects firmware setup access. With an administrator/setup password configured, and with permitted boot devices and boot order locked down, a user cannot enter firmware configuration to enable external media, change the startup order, or select an unapproved boot source. This prevents the common physical-console route for launching an OS installer from USB or optical media. Firmware security guidance also recommends controlling boot configuration and protecting setup access; see Microsoft: Windows security book. Together, these controls protect both firmware-level boot selection and bootloader-level startup modification.

QUESTION NO: 32

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:

SD card is the correct answer in the context of this Server+ question. A Type 1 hypervisor has a relatively small boot footprint and runs directly on the server hardware, so it is commonly deployed on dedicated boot media rather than on the local disks or network storage intended for virtual-machine datastores. An internal SD card provides a dedicated location for the hypervisor operating environment while leaving SAS and SATA storage available for virtual-machine files, application data, RAID volumes, or other server workloads. This separation also makes replacement or reinstallation of the hypervisor more straightforward because the boot medium is distinct from the primary data storage. Network-attached storage is generally used as shared storage for virtual machines and data, not as the preferred local boot target for the host. In production, the selected SD media must be enterprise-grade and supported by the hypervisor vendor, with appropriate endurance and redundancy considerations. Newer platform guidance may favor other persistent boot devices for some current hypervisor releases, but the dedicated SD-card installation model is the expected answer for the stated CompTIA Server+ scenario. See CompTIA's exam objectives resources and Broadcom's guidance on ESXi boot-media considerations.

QUESTION NO: 33

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** provide complementary layers of power redundancy for a server rack. Separate circuits establish independent upstream power paths, ideally sourced from distinct breakers and electrical distribution paths. Equipment with dual power supplies can then connect one supply to each circuit, so a failure involving one branch circuit does not necessarily interrupt server operation. This arrangement removes a common single point of failure in the rack's normal utility-power feed.

Multiple UPSs extend that design by supplying independent battery-backed power paths. A typical resilient implementation uses two appropriately sized UPS units, each supporting a separate circuit and power distribution unit, with dual-corded servers connected across both paths. If a UPS, its battery system, or its associated circuit requires maintenance or experiences a fault, the alternate path can continue delivering conditioned power to the server. The UPS units also provide short-term battery runtime during an outage, allowing time for standby generation to start or for an orderly shutdown if necessary. UPS redundancy should be designed with sufficient capacity for the intended load and with properly maintained batteries. See the CISA UPS guidance and APC guidance on redundant UPS systems.

QUESTION NO: 34

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:

Signal blocking is appropriate because it reduces the ability of radio-frequency and other electromagnetic emissions to pass through a facility's perimeter. In a datacenter, this can include appropriately designed electromagnetic shielding around sensitive areas, shielded cabling, filtered penetrations, and other emission-security controls. These measures help limit the unintentional external detection or interception of signals associated with systems and communications. The U.S. National Security Agency describes TEMPEST as a program concerned with compromising emanations and associated countermeasures, which supports using shielding and emission-control techniques where such exposure is a concern. See the NSA privacy-protection resources.

Reflective glass supports perimeter security by limiting outside observation through windows and reducing the visibility of interior activity, equipment, displays, and indicators. Window treatments with reflective or specialized films can provide a visual-observation barrier while maintaining the practical use of exterior windows. Used alongside signal-blocking measures, reflective glass addresses the physical exposure created by externally observable building activity and helps make datacenter operations less detectable from outside the facility. NIST's physical and environmental protection guidance recognizes the need to control physical access and protect information systems through facility safeguards; see NIST SP 800-53, Physical and Environmental Protection controls.

QUESTION NO: 35

A server is performing slowly, and users are reporting issues connecting to the application on that server. Upon investigation, the server administrator notices several unauthorized services 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: B C

Explanation:

The server is infected with a virus is a likely cause because malware can install or launch unapproved processes and services, consume processor, memory, storage, or network resources, and establish outbound communications to command-and-control infrastructure. Those effects directly align with both the degraded server/application performance and the discovery of unauthorized services communicating externally. NIST describes malware as software that compromises confidentiality, integrity, or availability, including software that can execute unwanted actions on a system. See NIST's malware definition.

The firewall rule for the server is misconfigured is also a likely contributing cause. A permissive or incorrectly scoped egress rule can allow unauthorized processes to communicate with an external host when that traffic should have been restricted. Server firewall policy should use least privilege, permitting only the specific inbound and outbound ports, protocols, destinations, and applications required for legitimate business services. Reviewing outbound rules, connection logs, and the external destination is therefore an appropriate part of containment and remediation. NIST guidance on firewall policy and rule configuration is available in SP 800-41 Rev. 1, Guidelines on Firewalls and Firewall Policy.

QUESTION NO: 36

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:

The required well-known ports are **21** for FTP, **22** for SSH, and **389** for LDAP. FTP uses TCP port 21 for its control connection, through which a client authenticates and sends commands. SSH uses TCP port 22 to provide encrypted remote command-line administration and secure file-transfer services such as SCP and SFTP. LDAP uses TCP and UDP port 389 for standard directory-service queries and related directory access; a firewall rule should use the protocol appropriate to the server's required LDAP traffic. These assignments are registered service names and port numbers, making them the expected ports to permit when configuring host-based firewall access for the stated services. In a production FTP deployment, administrators may additionally need to permit a defined passive-mode data-port range because FTP data transfers use a separate connection. That operational consideration does not change the requested standard service port. The Internet Assigned Numbers Authority lists FTP on port 21, SSH on port 22, and LDAP on port 389 in its Service Name and Transport Protocol Port Number Registry. Microsoft's Active Directory documentation also identifies port 389 as the standard LDAP port. See the IANA Service Name and Port Number Registry and Microsoft documentation on Active Directory firewall ports.

QUESTION NO: 37

A technician is investigating a server that unexpectedly restarted overnight. Which of the following sources should the technician review FIRST to determine the cause of the restart? (Choose two.)

A. System event logs

B. The server management controller logs

C. Web browser history

D. User home-directory contents

E. Printer queue history

ANSWER: A B

Explanation:

System event logs and **the server management controller logs** are the most appropriate initial sources. Operating-system system logs can record events related to unexpected shutdowns, stop errors, service failures, update installation, and restart activity. Reviewing the events immediately before the restart can help determine whether the operating system initiated the restart or encountered an error.

The server management controller, such as an iDRAC, iLO, or similar baseboard management controller, maintains hardware-level event information independently of the operating system. Its logs can identify power loss, thermal events, fan failures, power-supply faults, watchdog resets, and other hardware conditions that may not be recorded by the operating system. A web browser history and user home-directory contents are not reliable sources for diagnosing an unexpected server restart. Microsoft documents Windows event-log management at wevtutil documentation.

QUESTION NO: 38

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:

Perpetual is correct because a perpetual software license is purchased once and grants the organization the right to use the specified version of the software indefinitely. This licensing model aligns with the request to treat the acquisition as a capital expense during the current year, rather than as an ongoing operating expense that must be renewed periodically. Although optional support, maintenance, or upgrade coverage may be purchased separately, those services do not change the underlying perpetual right to continue using the licensed software version after the initial purchase. In practical server environments, this model can be appropriate when the organization wants a predictable up-front acquisition cost and does not require access to continuously released features through a recurring term agreement. Licensing is a Server+ objective area because administrators must be able to recognize the operational and financial implications of common licensing models when deploying and managing server software. CompTIA's Server+ exam page identifies software licensing concepts as part of the skills assessed: CompTIA Server+. Microsoft also distinguishes perpetual licensing from subscription licensing in its licensing guidance: Microsoft licensing plan options.

QUESTION NO: 39

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:

Confirm the network has a helper address for the DHCP server. DHCP clients initially send a DHCPDISCOVER message as a local broadcast because they do not yet have an IP address or a configured DHCP server. Broadcast traffic is normally contained within its originating VLAN, so a client in the newly created VLAN cannot directly reach a DHCP server located in another subnet. A DHCP relay, commonly configured as an `ip helper-address` on the VLAN's Layer 3 switch virtual interface or default-gateway router interface, receives that broadcast and forwards it as a unicast request to the DHCP server cluster. The relay also supplies gateway-address information, enabling the DHCP server to select the scope that

corresponds to the new VLAN's subnet. Since the scope has already been created and activated on an existing DHCP deployment, confirming that the new VLAN has a correctly configured relay target is the most direct first troubleshooting step. The relay address should point to a reachable DHCP server or an appropriate DHCP relay/load-balancing endpoint, and routing and any relevant firewall rules must permit the DHCP relay exchange. Microsoft documents DHCP relay-agent behavior and the importance of relaying requests across subnet boundaries in its DHCP overview. Cisco also describes configuring an IP helper address for DHCP relay.

QUESTION NO: 40

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:

Install the latest patches is a fundamental application-hardening practice. Application hardening reduces exposure to vulnerabilities by ensuring software is maintained at a supported and secure patch level. Vendors release patches to correct known security flaws, remediate defects that could enable unauthorized access or privilege escalation, and address vulnerabilities identified through ongoing security research. On a newly deployed server, the administrator should apply operating-system updates as well as updates for every installed server application, framework, runtime, library, and supporting component. Establishing an ongoing patch-management process is equally important: assess updates, test them when appropriate, deploy them on a defined schedule, and prioritize security updates for actively exploited or critical vulnerabilities. This minimizes the time that known weaknesses remain available to attackers and is an essential part of maintaining a hardened application environment. Guidance from the CISA Known Exploited Vulnerabilities Catalog emphasizes timely remediation of vulnerabilities known to be exploited, while the NIST Guide to Enterprise Patch Management Planning describes patch management as a core organizational process for reducing vulnerability risk.

QUESTION NO: 41

A server administrator is investigating intermittent application failures. The server event log shows repeated disk-related warnings, and the RAID controller reports that a physical drive is predictive failure. Which of the following actions should the administrator take? (Choose two.)

A. Verify that current backups are available

B. Replace the failing drive

C. Format the logical volume

D. Disable RAID monitoring

E. Initialize all disks in the array

F. Increase the page file size

ANSWER: A B

Explanation:

Verify that current backups are available should be performed because a predictive drive failure can progress to a complete failure and may expose other storage or recovery problems. Confirming that valid, recoverable backups exist protects the organization if the array becomes unavailable or if an error occurs during replacement and rebuild operations.

Replace the failing drive is also appropriate. Predictive-failure alerts indicate that the controller or drive has detected conditions associated with a likely future failure, such as excessive media errors or SMART-related indicators. In a fault-tolerant array with hot-swappable drive bays, the administrator should use a compatible replacement drive and follow the vendor-approved procedure so the controller can rebuild or restore redundancy. The array should then be monitored until it returns to a healthy state. NIST recommends maintaining backups to support recovery, and storage vendors identify predictive-failure status as an indicator that a drive should be replaced. See NIST SP 800-34 Rev. 1 and Dell PERC error codes and messages.

QUESTION NO: 42

A systems administrator is configuring an NTP service for servers in an Active Directory environment. Which of the following are benefits of maintaining accurate time synchronization? (Choose two.)

A. Kerberos authentication

B. Accurate event correlation

C. Increased disk capacity

D. Reduced CPU clock speed

E. Automatic RAID rebuilding

F. Enhanced VLAN tagging

ANSWER: A B

Explanation:

Kerberos authentication depends on time synchronization because Kerberos tickets include timestamps and have defined validity periods. If a client, server, or domain controller has excessive clock skew, ticket validation can fail and users can be unable to authenticate to domain resources. Accurate time across domain members helps ensure that authentication requests are accepted within the allowed time tolerance.

Accurate event correlation is also a benefit. Security, application, and operating-system logs are far more useful during troubleshooting and incident response when events from different systems have consistent timestamps. Administrators can accurately establish event order, identify related activity, and investigate incidents across servers, network devices, and security tools. Microsoft documents the importance of Windows Time for Kerberos authentication, and NIST identifies synchronized timestamps as important for audit-record correlation. See Microsoft Windows Time Service documentation and NIST SP 800-92: Guide to Computer Security Log Management.

QUESTION NO: 43

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 correct because they provide a practical, cost-effective possession-based authentication method for an MFA deployment. After a user enters their normal account credentials, an authenticator application on the user's enrolled mobile device receives a request that the user must approve. The mobile device serves as something the user has, while the

password is something the user knows; these are separate authentication factors. This approach can be deployed through an identity provider or MFA service without purchasing, issuing, tracking, and replacing dedicated hardware for every server administrator. It also centrally supports enrollment, revocation, logging, and policy enforcement across many servers. For stronger protection, the push workflow should use number matching or a comparable verification step so that the user actively confirms the specific sign-in attempt, rather than blindly approving a prompt. NIST's digital identity guidance describes MFA as using authentication factors from different categories and addresses out-of-band authentication methods: NIST SP 800-63B. CISA also recommends MFA and identifies push-based authenticators as a commonly available implementation method: CISA: Turn On MFA.

QUESTION NO: 44

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. Type 1

E. Type2

F. Guest

ANSWER: A D

Explanation:

The applicable selections are **Host** and **Type 1**. Interpreting the requirement as deploying virtual machines on a physical, bare-metal server, the host is the physical server platform that supplies the compute, memory, storage, and network resources on which virtualization is provided. A bare-metal host is therefore the foundational infrastructure required to run the virtualized workload. A Type 1 hypervisor is installed directly on that physical host's hardware rather than being installed as an application within a conventional desktop or server operating system. It manages hardware resources and presents virtual hardware to guest operating systems, enabling the virtual OS to run as a virtual machine. This architecture is normally selected for server virtualization because it provides direct hardware-level virtualization and centralized management of virtual machines. VMware describes ESXi as a bare-metal hypervisor installed directly on a physical server, and Microsoft explains that the Hyper-V hypervisor sits between hardware and operating systems to manage virtualized resources. See VMware ESXi Installation documentation and Microsoft Hyper-V overview.

QUESTION NO: 45 - (SIMULATION)

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.

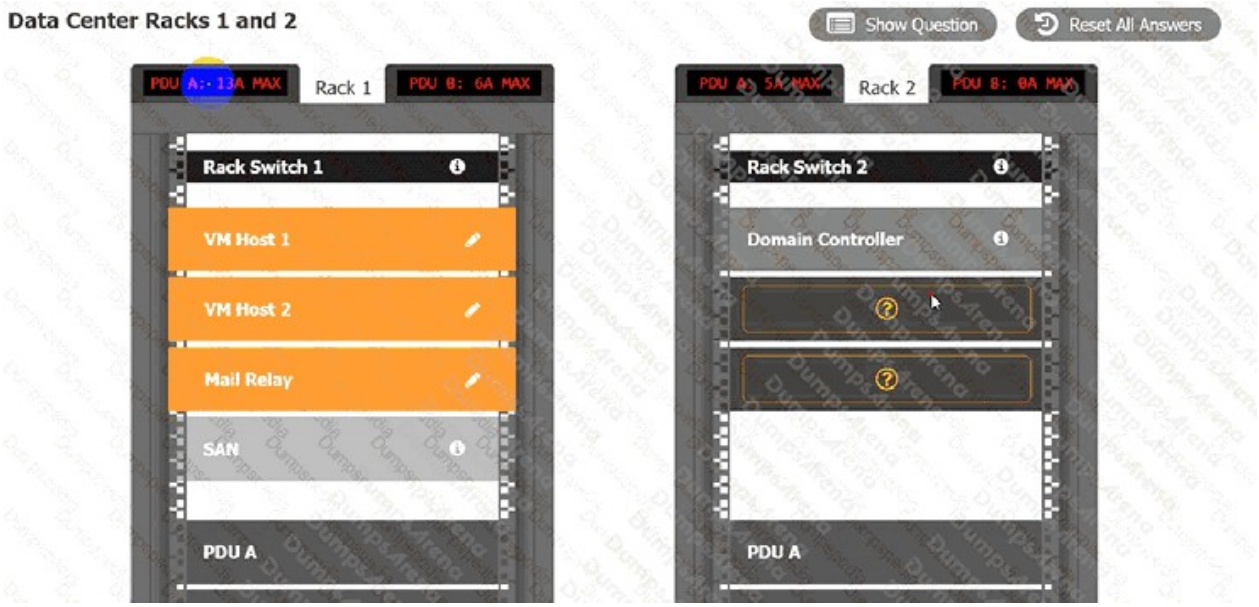
INSTRUCTIONS

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



ANSWER: See the explanation for the answer

Explanation:

Move VM Host 1 to Rack 2 and configure its PDU selection as PDU A.

VM Host 1 draws 4A. Moving it from Rack 1 PDU A to Rack 2 PDU A resolves the overload while preserving A/B power-path redundancy in both racks.

All PDUs are therefore at or below the required **9.6A (80%)** maximum, and each rack has equipment connected through both PDU A and PDU B.

Use the pencil icon on VM Host 1 to select `PDU A`, then place VM Host 1 into one of the open Rack 2 slots.

QUESTION NO: 46

Which of the following often-overlooked parts of the asset life cycle can cause the greatest number of issues in relation to PII exposure?

A. Usage

B. End-of-life

C. Procurement

D. Disposal

ANSWER: D

Explanation:

Disposal is correct because retired servers, storage devices, removable media, and other IT assets can retain personally identifiable information (PII) long after they leave active service. Data may persist in user files, databases, logs, swap space, RAID members, backup partitions, and device configuration areas. Simply deleting files or reformatting a drive does not necessarily make that information unrecoverable. If media are discarded, resold, returned under lease, sent for repair, or transferred to another party without verified sanitization, the organization can lose control of sensitive information and create a reportable privacy or compliance incident.

A sound disposal process therefore includes maintaining asset custody records, identifying the data classification of the asset, selecting a sanitization method appropriate to the media type and reuse requirements, and documenting completion. NIST SP 800-88 Rev. 1 describes media sanitization approaches, including clear, purge, and destroy methods, and stresses

verification and documentation as part of a defensible process. CompTIA Server+ objectives also emphasize secure decommissioning and disposal practices for server hardware and storage. These controls reduce the chance that residual PII is exposed when an asset reaches the end of its useful life. See NIST SP 800-88 Rev. 1 and CompTIA exam objectives.

QUESTION NO: 47

A server administrator is installing environmental monitoring for a rack of production servers. Which of the following conditions should the administrator monitor to prevent hardware damage and service interruptions? (Choose two.)

A. Temperature

B. Humidity

C. MAC address table size

D. DNS zone transfers

E. User login count

ANSWER: A B

Explanation:

Temperature and **humidity** are critical environmental conditions to monitor in a server rack. Excessive temperature can cause processors, memory, storage devices, and power supplies to overheat, potentially resulting in thermal throttling, unexpected shutdowns, or component failure. Monitoring temperature allows administrators to identify cooling failures, blocked airflow, and abnormal heat conditions before availability is affected.

Humidity must also remain within the equipment manufacturer's supported operating range. Low humidity can increase the likelihood of electrostatic discharge, while excessive humidity can cause condensation and corrosion. Environmental sensors and alert thresholds enable administrators to respond to cooling or humidity issues before server hardware is damaged. ASHRAE provides environmental guidance for data-processing environments in its Datacom Series, and NIST includes environmental monitoring considerations in SP 800-53.

QUESTION NO: 48

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:

Checking the hosts file and checking the DNS server are the appropriate next steps because successful access by IP address establishes that basic Layer 3 network connectivity to the new server is working. The remaining symptom is specifically name resolution: a client is attempting to use the server’s name, but that name is not being translated into the correct IP address. A local hosts file can provide a static hostname-to-address mapping and is consulted as part of the operating system’s name-resolution process. The administrator should verify that no stale, incorrect, or missing local

mapping is affecting access. Microsoft documents hosts-file behavior in its DNS concepts documentation. The DNS server should also be checked to confirm that the new server has an appropriate forward lookup record, that the record points to the intended address, and that clients are configured to query the authoritative or otherwise correct DNS service. DNS is the service designed to resolve host names to IP addresses, so validating its records and availability directly addresses a failure that occurs only when the name is used. Microsoft also provides guidance for managing DNS records.

QUESTION NO: 49

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:

Disclose the incident. is an essential breach-response measure because affected parties, organizational leadership, regulators, law enforcement, customers, or partners may need timely notification. The exact notification obligation and deadline depend on the organization's incident-response plan, contractual commitments, the type and location of affected data, and applicable laws or regulations. Prompt, coordinated disclosure supports transparency, enables affected individuals to take protective action, and ensures communications are managed by authorized legal, privacy, and incident-response stakeholders.

Identify the exploited vulnerability. is also central to incident handling. After initial containment, responders must determine the attack vector, affected systems and accounts, scope of data exposure, and the security weakness the attacker used. This evidence-based investigation allows the organization to remediate the root cause through actions such as patching, configuration correction, credential rotation, or improved controls. Identifying the vulnerability also helps prevent recurrence and provides the technical facts needed for accurate reporting, recovery, and lessons learned. NIST's incident-handling guidance emphasizes analysis, containment, eradication, recovery, and post-incident activity; CISA likewise recommends investigating the intrusion and addressing the initial access path. See NIST SP 800-61 Rev. 2 and CISA Incident Response.

QUESTION NO: 50

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:

The site has partial infrastructure and some data.

A warm site is a disaster-recovery location that provides a meaningful portion of the resources needed to resume business operations, but it is not continuously ready to assume production workloads. It typically includes prearranged facilities, hardware, network capability, and selected systems or services. The organization must still complete configuration, restore or synchronize required data, and validate applications before the location can support normal operations. Because the environment contains some infrastructure and some recoverable data, recovery is substantially faster than building an environment from scratch, while still requiring more activation work than an immediately operational recovery environment.

This description aligns with the practical recovery-planning purpose of a warm site: it balances recovery speed and cost by maintaining only the components needed to accelerate restoration. The exact amount of infrastructure and data is determined by the organization's recovery time and recovery point objectives, so documented recovery procedures, current backups, and regular exercises are essential to make the warm site usable during an incident. NIST's definition and recovery-planning guidance discuss alternate processing sites and their readiness levels: NIST CSRC: warm site and NIST SP 800-34 Rev. 1.

QUESTION NO: 51

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:

The required ports are **21**, **22**, and **389**. SSH uses TCP port 22 as its well-known listening port, so permitting inbound TCP/22 allows remote Secure Shell administration connections to the server. FTP's control connection uses TCP port 21; this is the port a client initially contacts for authentication and FTP commands. LDAP directory-service traffic uses TCP and UDP port 389, and a host firewall rule allowing LDAP normally includes TCP/389 for directory queries and sessions. Because the question asks for the standard ports associated with SSH, FTP, and LDAP, these three well-known port numbers are the appropriate firewall configuration choices. In a production configuration, the administrator should also account for the organization's specific FTP data-channel mode and any requirement for encrypted alternatives such as LDAPS, but those details do not change the three standard service ports requested here. The Internet Assigned Numbers Authority lists these assignments in its Service Name and Transport Protocol Port Number Registry: IANA service names and port numbers. Microsoft's LDAP documentation also identifies port 389 as the default LDAP port: Configure a firewall for Active Directory domains and trusts.

QUESTION NO: 52

An administrator notices a server is offline. 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:

A DIMM has failed is the most likely cause. “Configuring Memory” occurs during the server’s preboot hardware initialization, when firmware performs memory discovery, training, configuration, and validation before the operating system begins loading. A failed or otherwise defective DIMM can prevent this initialization from completing, leaving the server stalled at this stage even after a reboot. The ability to access out-of-band remote management is consistent with this diagnosis: the management controller operates independently of the host operating system and can remain reachable while the host itself cannot complete POST and boot. A practical next step is to review the out-of-band event log and system-health information for memory-related alerts, such as an uncorrectable ECC error or a failed DIMM location. Following the server manufacturer’s maintenance guidance, the administrator can then power down the system, reseat the indicated module, and test or replace it with supported memory. CompTIA’s Server+ exam objectives include diagnosing server hardware faults, including memory-related failures. See the CompTIA Server+ certification page and Dell’s PowerEdge memory troubleshooting guidance.

QUESTION NO: 53

A systems administrator is reviewing capacity trends for a file server. Which of the following should the administrator collect to determine whether storage capacity must be expanded? (Choose two.)

A. Available disk space

B. Disk growth rate

C. CPU utilization

D. Network latency

E. Number of user accounts

ANSWER: A B

Explanation:

Available disk space and **disk growth rate** are the most useful metrics for storage-capacity planning. Available disk space identifies the current amount of free capacity on each filesystem or volume. A volume nearing capacity can prevent users and applications from creating files, writing logs, installing updates, or completing backups.

Disk growth rate shows how quickly the used capacity is increasing over time. By comparing historical utilization with the growth rate, the administrator can forecast when a volume will reach a defined threshold and schedule an expansion before service is affected. CPU utilization and network latency can be useful performance metrics, but they do not directly determine whether additional storage capacity is required. NIST describes capacity planning as an activity that uses utilization and trend information to maintain adequate information-system resources. See NIST SP 800-137.

QUESTION NO: 54

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:

Connect to the server from a crash cart. A crash cart provides direct, local console access to a physical server when normal remote-management paths are unavailable. It commonly includes a monitor, keyboard, mouse, and the appropriate display and USB cables, allowing the technician to attach directly to the server's video and peripheral ports. This gives the technician visibility of POST messages, firmware or BIOS/UEFI setup, boot errors, operating-system startup behavior, and any local diagnostic output that cannot be reached through the network.

In this situation, remote management software cannot establish a usable session because the out-of-band management interface has not been configured. A shared KVM path is also unavailable because the servers are not connected to a KVM switch. Using a crash cart is therefore the practical best-practice method to establish an immediate local console session without relying on existing network connectivity or prior management-controller configuration. Once access is restored, the technician can investigate the fault and, if appropriate, configure the out-of-band management controller for future remote access. CompTIA's Server+ exam objectives include managing servers through local and remote access methods and using appropriate troubleshooting tools. See the CompTIA exam objectives page and Dell's iDRAC configuration guidance.