

CompTIA.SK0-005.vOct-2023.by.Herry.87q

Number: SK0-005
Passing Score: 800
Time Limit: 120
File Version: 12.0

Exam Code: SK0-005
Exam Name: CompTIA Server+ Certification Exam



Exam A

QUESTION 1

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

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

Correct Answer: D

Section:

Explanation:

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

QUESTION 2

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

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

Correct Answer: D

Section:

Explanation:

The installation of an OS contained entirely within another OS installation is described as a guest. A guest is a term that refers to a virtual machine (VM) that runs on top of a host operating system (OS) using a hypervisor or a virtualization software. A guest can have a different OS than the host, and can run multiple applications or services independently from the host. A guest can also be isolated from the host and other guests for security or testing purposes.

QUESTION 3

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

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

Correct Answer: B

Section:

Explanation:

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

QUESTION 4

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

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

Correct Answer: A

Section:

Explanation:

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

Reference:

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

QUESTION 5

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

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

Correct Answer: D

Section:

Explanation:

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

QUESTION 6

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

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

Correct Answer: B

Section:

Explanation:

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

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

QUESTION 7

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

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

Correct Answer: D

Section:

Explanation:

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

QUESTION 8

A server administrator has noticed that the storage utilization on a file server is growing faster than planned. The administrator wants to ensure that, in the future, there is a more direct relationship between the number of users using the server and the amount of space that might be used. Which of the following would BEST enable this correlation?

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

Correct Answer: C

Section:

Explanation:

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

QUESTION 9

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

- A. Copy the fileshare to an LTO-4 tape drive
- B. Configure a new incremental backup job for the fileshare

- C. Create an additional partition and move a copy of the fileshare
- D. Create a snapshot of the fileshare

Correct Answer: D

Section:

Explanation:

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

QUESTION 10

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

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

Correct Answer: C

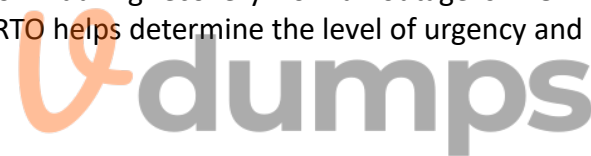
Section:

Explanation:

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

Reference:

<https://whatis.techtarget.com/definition/recovery-time-objective-RTO>



QUESTION 11

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

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

Correct Answer: C

Section:

Explanation:

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

QUESTION 12

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

Correct Answer: A, B

Section:

Explanation:

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

Reference:

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

QUESTION 13

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

># Enforcing

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

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



Correct Answer: A

Section:

Explanation:

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

Reference:

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

QUESTION 14

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

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

Correct Answer: B

Section:

Explanation:

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

QUESTION 15

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

- A. DNS server
- B. File server
- C. DHCP server
- D. Web server

Correct Answer: D

Section:

Explanation:

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

Reference:

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

QUESTION 16

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

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

Correct Answer: D

Section:

Explanation:

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

QUESTION 17

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

Correct Answer: D, E

Section:

Explanation:

The most likely reasons why the technician is unable to access a server's package repository internally or externally are that the external firewall is blocking access and that the default gateway is incorrect. A package repository is a source of software packages that can be installed or updated on a server using a package manager tool. A package repository can be accessed over a network using a URL or an IP address. However, if there are any network issues or misconfigurations, the access to the package repository can be blocked or failed. An external firewall is a device or software that controls the incoming and outgoing network traffic based on predefined rules or policies. An external firewall can block access to a package repository if it does not allow traffic on certain ports or protocols that are used by the package manager tool. A default gateway is a device or address that routes network traffic from one network to another network. A default gateway can be incorrect if it does not match the actual device or address that connects the server's network to other networks, such as the internet. An incorrect default gateway can prevent the server from reaching the package repository over other networks.

QUESTION 18

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

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

Correct Answer: A

Section:

Explanation:

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

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

QUESTION 19

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

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

Correct Answer: A

Section:

Explanation:

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

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

QUESTION 20

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

Correct Answer: C

Section:

Explanation:

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

Reference:

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

QUESTION 21

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

- A. Disable port 389 on the server
- B. Move traffic from port 389 to port 443
- C. Move traffic from port 389 to port 637
- D. Enable port 389 for web traffic



Correct Answer: A

Section:

Explanation:

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

QUESTION 22

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

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

Correct Answer: B

Section:

Explanation:

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

support can help raise awareness and reinforce best practices among the employees.

QUESTION 23

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

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

Correct Answer: C

Section:

Explanation:

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

QUESTION 24

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

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

Correct Answer: A

Section:

Explanation:

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

Reference:

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

QUESTION 25

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

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

Correct Answer: B



Section:

Explanation:

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

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

QUESTION 26

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

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

Correct Answer: C

Section:

Explanation:

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

QUESTION 27

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

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

Correct Answer: B

Section:

Explanation:

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

QUESTION 28

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

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

Correct Answer: A

Section:

Explanation:

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

QUESTION 29

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

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

Correct Answer: B

Section:

Explanation:

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

QUESTION 30

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

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

Correct Answer: C

Section:

Explanation:

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

QUESTION 31

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

\\server1\ITDept\

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

- A. eSATA

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

Correct Answer: C

Section:

Explanation:

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

QUESTION 32

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

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

Correct Answer: C

Section:

Explanation:

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

QUESTION 33

A technician is connecting a Linux server to a share on a NAS. Which of the following is the MOST appropriate native protocol to use for this task?

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

Correct Answer: D

Section:

Explanation:

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

QUESTION 34

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

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

Correct Answer: C

Section:

Explanation:

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

Reference:

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

QUESTION 35

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

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

Correct Answer: A

Section:

Explanation:

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

QUESTION 36

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

Correct Answer: D

Section:

Explanation:

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

QUESTION 37

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

- A. Perform a root cause analysis
- B. Develop a plan of action
- C. Document the findings
- D. Escalate the issue
- E. Scope the issue
- F. Notify the users

Correct Answer: C, F

Section:

Explanation:

The steps in the troubleshooting theory that should be performed after a solution has been implemented are document the findings and notify the users. The troubleshooting theory is a systematic process of identifying and resolving problems or issues with a system or device. The troubleshooting theory consists of several steps that can be summarized as follows:

Identify the problem: Gather information, scope the issue, establish a theory of probable cause. Establish a plan of action: Test the theory, determine next steps, escalate if necessary. Implement the solution: Execute the plan, verify functionality, prevent recurrence. Document the findings: Record actions taken, outcomes achieved, lessons learned. Notify the users: Communicate resolution status, confirm satisfaction, provide follow-up. Documenting the findings is an important step that helps create a record of what was done and why, what worked and what didn't, and what can be improved or avoided in the future. Documenting the findings can also help with reporting, auditing, compliance, or training purposes. Notifying the users is another important step that helps inform the affected parties of what was done and how it was resolved, confirm that the problem is fixed and that they are satisfied with the outcome, and provide any follow-up instructions or recommendations.

QUESTION 38

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

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

Correct Answer: A

Section:

Explanation:

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

QUESTION 39

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

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

D. perform a backup restore.

Correct Answer: B

Section:

Explanation:

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

QUESTION 40

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

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

Correct Answer: D

Section:

Explanation:

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

Reference: [https://docs.vmware.com/en/VMwarevSphere/](https://docs.vmware.com/en/VMwarevSphere/6.7/com.vmware.vsphere.vm_admin.doc/GUID-8254CD05-CC06-491D-BA56-A773A32A8130.html)

[6.7/com.vmware.vsphere.vm_admin.doc/GUID-8254CD05-CC06-491D-BA56-A773A32A8130.html](https://docs.vmware.com/en/VMwarevSphere/6.7/com.vmware.vsphere.vm_admin.doc/GUID-8254CD05-CC06-491D-BA56-A773A32A8130.html)

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

QUESTION 41

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

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

Correct Answer: A

Section:

Explanation:

Monitors Network Traffic

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

QUESTION 42

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

Correct Answer: C, E

Section:

Explanation:

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

QUESTION 43

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

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

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

IP address: 192.168.1.1/24

Default gateway: 192.168.10.1

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

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

Correct Answer: A

Section:

Explanation:

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

QUESTION 44

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

administrator MOST likely need to install?

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

Correct Answer: D

Section:

Explanation:

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

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

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

QUESTION 45

Users at a company work with highly sensitive dat

- A. The security department implemented an administrative and technical control to enforce leastprivilege access assigned to files. However, the security department has discovered unauthorized data exfiltration. Which of the following is the BEST way to protect the data from leaking?
- B. Utilize privacy screens.
- C. Implement disk quotas.
- D. Install a DLP solution.
- E. Enforce the lock-screen feature.



Correct Answer: C

Section:

Explanation:

Components of a Data Loss Solution

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

QUESTION 46

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

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

Correct Answer: C

Section:

Explanation:

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

QUESTION 47

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

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

Correct Answer: B

Section:

Explanation:

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

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

QUESTION 48

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

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



Correct Answer: B

Section:

Explanation:

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

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

QUESTION 49

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

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

Correct Answer: A

Section:

Explanation:

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

QUESTION 50**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.

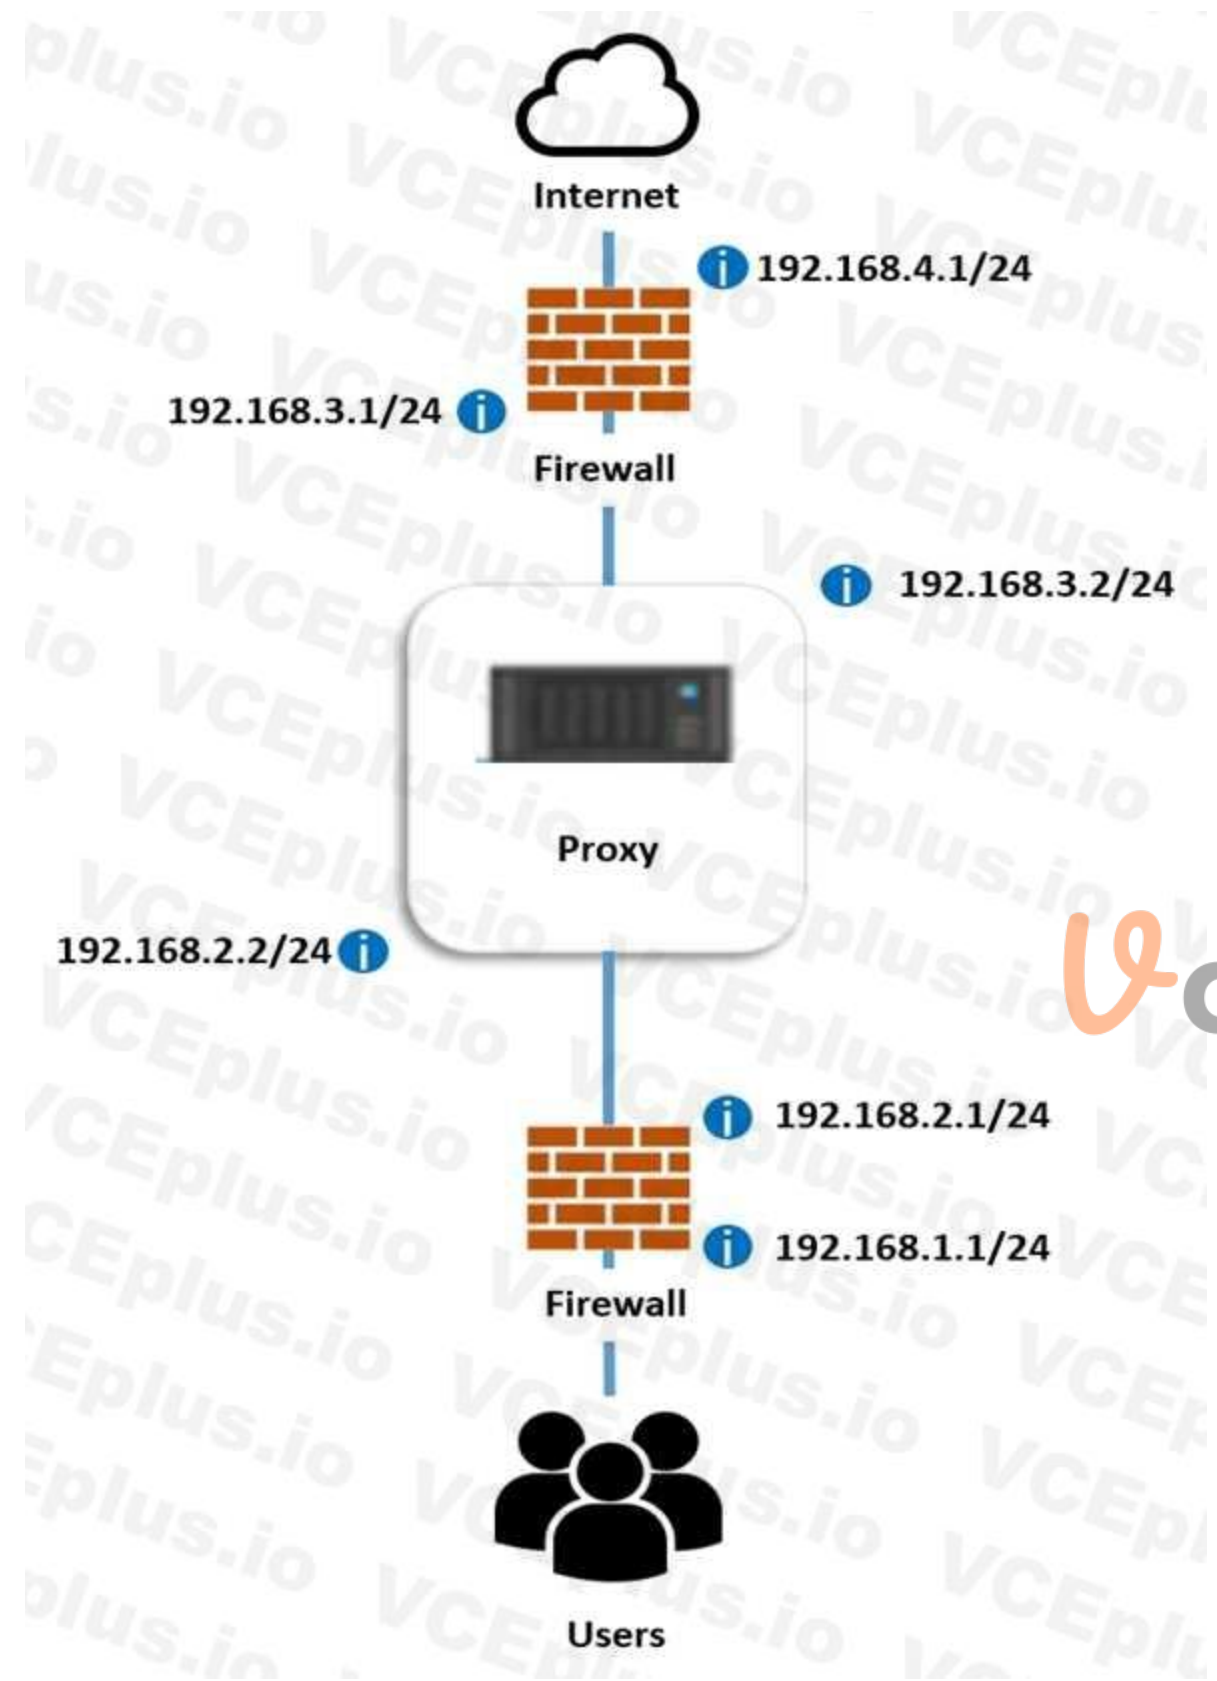
INSTRUCTIONS

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.





 **Vdumps**

Proxy Server Routing Table

Destination	Netmask	Gateway	Interface
0.0.0.0	0.0.0.0	<ul style="list-style-type: none"> 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 	<ul style="list-style-type: none"> 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	<ul style="list-style-type: none"> 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 	<ul style="list-style-type: none"> 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

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 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 0.0.0.0 192.168.3.1 255.255.255.0 192.168.4.0 192.168.2.1 192.168.2.0

Section:

Explanation:

QUESTION 51

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:

- A. The OS must be installed in a separate disk partition. In case of hard drive failure, it cannot be affected.
- B. Application data IOPS performance is a must.
- C. 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.)
- D. Install the OS on a RAID 0 array.
- E. Install the OS on a RAID 1 array.
- F. Configure RAID 1 for the application data.

- G. Configure RAID 5 for the application data.
- H. Use SSD hard drives for the data application array.
- I. Use SATA hard drives for the data application array.
- J. Use a single JBOD for OS and application data.

Correct Answer: B, D, E

Section:

Explanation:

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

Reference:

<https://phoenixnap.com/kb/raid-levels-and-types>

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

QUESTION 52

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

```
ipconfig /all

IPv4 address: 192.168.1.5
Subnet mask: 255.255.255.0
Default gateway: 192.168.1.1

pinging 192.168.1.1 with 32 bytes of data:

Request timed out
Reply from 192.168.1.1: bytes=32 time<1ms TTL=128
Request timed out
Reply from 192.168.1.1: bytes=32 time<1ms TTL=128
```

Which of the following caused the issue?

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

Correct Answer: A

Section:

Explanation:

The ping command output shows that the NIC has an IP address of 192.168.1.100 and a default gateway of 192.168.1.1. However, when the technician tries to ping the default gateway, the reply comes from

another IP address: 192.168.1.101. This means that there is another device on the network that has the same IP address as the default gateway, and it is responding to the ping request instead of the intended destination.

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

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

Reference:

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

QUESTION 53

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

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

Correct Answer: C

Section:

Explanation:

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

Reference: <https://pcgearhead.com/installing-a-new-gpu/>

QUESTION 54

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

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

Correct Answer: D

Section:

Explanation:

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

<https://www.ibm.com/cloud/learn/service-level-agreements>

QUESTION 55

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

- A. RTO
- B. MTBF
- C. PRO

D. MTTR

Correct Answer: A

Section:

Explanation:

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

QUESTION 56

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

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

Correct Answer: D

Section:

Explanation:

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

QUESTION 57

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

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

Correct Answer: C

Section:

Explanation:

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

QUESTION 58

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

- A. Restore the server from a snapshot.
- B. Restore the server from backup.

- C. Swap the drive and initialize the disk.
- D. Swap the drive and initialize the array.

Correct Answer: C

Section:

Explanation:

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

QUESTION 59

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

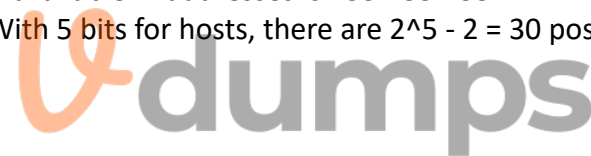
- A. 255.255.255.0
- B. 255.255.255.128
- C. 255.255.255.224
- D. 255.255.255.252

Correct Answer: C

Section:

Explanation:

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



QUESTION 60

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

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

Correct Answer: A

Section:

Explanation:

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

QUESTION 61

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

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

Correct Answer: A

Section:

Explanation:

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

QUESTION 62

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



Correct Answer: C, D

Section:

Explanation:

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

QUESTION 63

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

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

Correct Answer: B

Section:

Explanation:

To disable root login over SSH, the server administrator needs to edit the SSH configuration file located at /etc/ssh/sshd_config. This file contains various settings for the SSH daemon that runs on the server and accepts incoming SSH connections. The administrator needs to find the line that says PermitRootLogin and change it to no or comment it out with a # symbol. Then, the administrator needs to restart the SSH service for the changes to take effect. Reference:

<https://www.howtogeek.com/828538/how-and-why-to-disable-root-login-over-ssh-on-linux/>

QUESTION 64

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

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

Correct Answer: A

Section:

Explanation:

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

<https://www.techrepublic.com/article/understanding-raid-controller-write-cache/>

QUESTION 65

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

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



Correct Answer: B

Section:

Explanation:

Log rotation is a process that periodically renames, compresses, and deletes old log files to free up disk space and keep log files manageable. Log rotation can be configured using tools such as logrotate or cron on Linux systems, or using Windows Task Scheduler or PowerShell scripts on Windows systems. Log rotation can also help with log analysis and troubleshooting by making it easier to find relevant information in smaller and more recent log files. Reference:

<https://www.mezmo.com/learn-log-management/what-is-log-rotation-how-does-it-work> <https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/logman>

QUESTION 66

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

Correct Answer: A, C

Section:

Explanation:

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

<https://www.geeksforgeeks.org/what-is-hot-swapping/> <https://www.howtogeek.com/268249/whatis-hot-swapping-and-what-devices-support-it/>

QUESTION 67

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

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

Correct Answer: B

Section:

Explanation:

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

<https://www.pcmag.com/encyclopedia/term/concurrent-use-license> <https://www.techopedia.com/definition/1440/software-licensing>

QUESTION 68

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

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

Correct Answer: B

Section:

Explanation:

If users cannot access a new server by name, but the server does respond to a ping request using its IP address, it means that there is a problem with name resolution. The DNS (Domain Name System) is a service that maps hostnames to IP addresses and vice versa. Therefore, the best step to perform next is to examine the DNS to see if the new server record exists and matches its IP address. If not, the DNS record needs to be added or updated accordingly. Running the tracert command from a workstation would not help with name resolution, as it only shows the route taken by packets to reach a destination by IP address.

Correcting the missing DHCP scope would not help either, as DHCP (Dynamic Host Configuration Protocol) only assigns IP addresses and other network settings to clients, but does not resolve names. Updating the workstation hosts file would be a temporary workaround, but not a permanent solution, as it would require manually editing every workstation's hosts file with the new server's name and IP address. Reference:

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

QUESTION 69

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

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

Correct Answer: D

Section:

Explanation:

Teaming is a technique that can be configured on a server for network redundancy. Teaming involves combining two or more network adapters into a single logical unit that acts as one network interface. This way, if one network adapter fails, another one can take over without disrupting network connectivity. Teaming can also improve network performance by load balancing traffic across multiple network adapters.

Clustering is a technique that involves grouping two or more servers together to act as one system for high availability and fault tolerance. Virtualizing is a technique that involves creating multiple virtual machines on a single physical server to optimize resource utilization and flexibility. Cloning is a technique that involves creating an exact copy of a server's configuration and data for backup or migration purposes. Reference:

<https://docs.microsoft.com/en-us/windowsserver/networking/technologies/nic-teaming/nic-teaming>

<https://www.techopedia.com/definition/19588/clustering>

<https://www.techopedia.com/definition/4790/virtualization>

<https://www.techopedia.com/definition/4776/cloning>

QUESTION 70

An administrator is investigating a physical server that will not boot into the OS. The server has three hard drives configured in a RAID 5 array. The server passes POST, but the OS does not load. The administrator verifies the CPU and RAM are both seated correctly and checks the dual power supplies. The administrator then verifies all the BIOS settings are correct and connects a bootable USB drive in the server, and the OS loads correctly. Which of the following is causing the issue?

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



Correct Answer: C

Section:

Explanation:

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

<https://www.technewstoday.com/what-is-a-raid-5/>

QUESTION 71

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

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

Correct Answer: C

Section:

Explanation:

RTO (Recovery Time Objective) is a measure of how much downtime an organization can tolerate during an unplanned outage. It is the maximum time allowed for restoring normal operations after a disaster. RTO is one of the key metrics for disaster recovery planning and testing. SLA (Service Level Agreement) is a contract that defines the expected level of service and performance between a provider and a customer. BIA (Business Impact Analysis) is a process that identifies and evaluates the potential effects of a disaster on critical business functions and processes. MTTR (Mean Time To Repair) is a measure of how long it takes to fix a failed component or system. Reference:

<https://parachute.cloud/rto-vs-rpo/> <https://www.techopedia.com/definition/13622/service-levelagreement-sla> <https://www.techopedia.com/definition/1032/business-impact-analysis-bia>

<https://www.techopedia.com/definition/8239/mean-time-to-repair-mttr>

QUESTION 72

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

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

Correct Answer: A

Section:

Explanation:

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

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

QUESTION 73

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

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

Correct Answer: B

Section:

Explanation:

A hot site is a type of disaster recovery site that has all the equipment and data ready to resume operations as soon as possible after a disaster. A hot site is usually located in a different geographic area than the primary site and has redundant power, cooling, network, and security systems. A hot site is best for the company that wants to recover in the shortest possible time and does not want customer data to egress to non-company equipment. A warm site is a type of disaster recovery site that has some equipment and data ready, but requires some configuration and restoration before resuming operations. A cold site is a type of disaster recovery site that has only basic infrastructure and space available, but requires significant setup and installation before resuming operations. Cloud recovery is a type of disaster recovery service that uses cloud-based resources and platforms to store backups and restore data and applications after a disaster. Reference:

<https://www.techopedia.com/definition/11172/hot-site>

<https://www.techopedia.com/definition/11173/warm-site>

<https://www.techopedia.com/definition/11174/cold-site>

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

QUESTION 74

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

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

Correct Answer: D

Section:

Explanation:

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

QUESTION 75

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

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



Correct Answer: A

Section:

Explanation:

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

Reference:

<https://www.techopedia.com/definition/16293/mantrap>

<https://www.techopedia.com/definition/1437/bollard>

<https://www.techopedia.com/definition/23961/geofencing>

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

QUESTION 76

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

- A. A local firewall is blocking authentication requests.
- B. The users have expired passwords

- C. The system clock is off by more than five minutes
- D. The server has no access to the LDAP host

Correct Answer: C

Section:

Explanation:

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

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

QUESTION 77

Which of the following BEST describes a warm site?

The site has all infrastructure and live data.

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

Correct Answer: D

Section:

Explanation:

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

<https://www.enterprisestorageforum.com/management/disaster-recovery-site/> <https://www.techopedia.com/definition/3780/warm-site>

QUESTION 78

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

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

Correct Answer: D

Section:

Explanation:

RAID 10 is the RAID level that will yield the best combination of available space and redundancy when configuring a new server for use as a database server with two mirrored drives for the operating system and three drive bays remaining for storage. RAID 10, also known as RAID 1+0, is a RAID configuration that combines disk mirroring and disk striping to protect data. It requires a minimum of four disks and stripes data across mirrored pairs. As long as one disk in each mirrored pair is functional, data can be retrieved. RAID 10 provides high performance, fault tolerance, and fast recovery, but it reduces storage capacity by half. RAID 0 is a RAID configuration that splits data across two or more drives without parity or redundancy. It improves performance but offers no fault tolerance. If one drive fails in RAID 0, all data is lost and the system cannot boot. RAID 1 is a RAID configuration that duplicates data across two or more drives. It provides fault tolerance and improves read performance, but reduces storage capacity by half. If one drive fails in RAID 1, the other drive can continue to operate without data loss or system downtime. RAID 5 is a RAID configuration that stripes data across three or more drives with parity information. It provides fault

tolerance and improves performance, but reduces storage capacity by one drive's worth of space. RAID 5 can tolerate one drive failure without data loss, but not two or more. Reference: <https://www.howtogeek.com/199068/how-to-upgrade-your-existing-hard-drive-in-under-an-hour/> <https://www.howtogeek.com/202794/what-is-the-difference-between-127.0.0.1-and-0.0.0.0/>

QUESTION 79

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

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

Correct Answer: D

Section:

Explanation:

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

<https://www.howtogeek.com/428483/what-is-end-to-end-encryption-and-why-does-it-matter/> <https://www.howtogeek.com/443611/how-to-encrypt-your-macs-system-drive-removable-devices-and-individual-files/> <https://www.howtogeek.com/199068/how-to-upgrade-your-existing-harddrive-in-under-an-hour/> <https://www.howtogeek.com/202794/what-is-the-difference-between-127>

QUESTION 80

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

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

Correct Answer: A

Section:

Explanation:

A tabletop exercise is a type of disaster recovery testing exercise that is primarily used to discuss incident response strategies for critical systems without affecting production data. A tabletop exercise is a discussion-based session where team members meet in an informal, classroom setting to review their roles and responsibilities during an emergency and their responses to a hypothetical scenario. A facilitator guides the participants through the discussion and evaluates the strengths and weaknesses of the preparedness program. A tabletop exercise does not involve any actual deployment of resources or activation of systems. A backup recovery test (B) is a type of disaster recovery testing exercise that involves restoring data from backup media to verify its integrity and availability. A backup recovery test may affect production data if it is not performed on a separate environment. A live failover (C) is a type of disaster recovery testing exercise that involves switching operations from a primary site to a secondary site in case of a failure or disruption. A live failover may affect production data if it is not performed on a simulated environment. A hot-site visit audit (D) is a type of disaster recovery testing exercise that involves inspecting and evaluating a hot site, which is a backup location that has fully operational equipment and resources to resume business operations in case of a disaster. A hot-site visit audit does not involve any discussion of incident response strategies or simulation of scenarios. Reference: 1 <https://www.ready.gov/testingexercises> 2 <https://www.ready.gov/exercises>

QUESTION 81

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

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

Correct Answer: D

Section:

Explanation:

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

QUESTION 82

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

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



Correct Answer: B

Section:

Explanation:

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

QUESTION 83

A software developer is unable to reach an internal website. The developer's attempt to ping the FQDN returns the following IP address: 104.18.17.32. Which of the following is the most likely reason for this result?

- A. The NIC is set to DHCP.
- B. The default gateway is misconfigured.
- C. The primary DNS server is 8.8.8.8.
- D. There is a manual entry in the hosts file.

Correct Answer: D

Section:

Explanation:

The most likely reason for this result is that there is a manual entry in the hosts file that maps the FQDN to an incorrect IP address (104.18.17.32). The hosts file is a text file that contains mappings of hostnames or domain names to IP addresses, which are used by the operating system to resolve names before querying DNS servers on the network or internet. The hosts file can be used to override DNS settings or block access to certain websites by redirecting them to different IP addresses, such as localhost (127.0.0.1) or invalid addresses (0.0.0.0). If there is a manual entry in the hosts file that conflicts with DNS records, it can cause name resolution errors or connectivity issues. Reference: [CompTIA Server+ Certification Exam Objectives], Domain 4.0: Networking, Objective 4.1: Given a scenario, configure network settings for servers.

QUESTION 84

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

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

Correct Answer: D

Section:

Explanation:

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



QUESTION 85

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.

Correct Answer: A

Section:

Explanation:

The most likely issue that prevents the newly added server from seeing any of the LUNs on the SAN is that a single-mode fiber cable is used in place of multimode. A single-mode fiber cable is a type of optical fiber cable that has a small core diameter and allows only one mode of light to propagate through it. A single-mode fiber cable can transmit data over long distances at high speeds, but it requires more expensive transceivers and connectors than multimode fiber cables. A multimode fiber cable is a type of optical fiber cable that has a larger core diameter and allows multiple modes of light to propagate through it. A multimode fiber cable can transmit data over short distances at lower speeds than single-mode fiber cables, but it is more compatible and cost-effective than singlemode fiber cables. If a single-mode fiber cable is used in place of multimode, it can cause signal loss, attenuation, or mismatch between the devices. Reference: [CompTIA Server+ Certification Exam Objectives], Domain 3.0: Storage, Objective 3.2: Given a scenario, compare and contrast various storage technologies.

QUESTION 86

A server administrator is currently working on an incident. Which of the following steps should the administrator perform before resolving the issue?

- A. Inform the impacted users.
- B. Make the changes to the system.
- C. Determine the probable causes.
- D. Identify changes to the server.

Correct Answer: C

Section:

Explanation:

The step that the server administrator should perform before resolving the issue is to determine the probable causes. This step is part of the troubleshooting process that follows a logical and systematic approach to identify and solve problems with servers and applications. The troubleshooting process consists of several steps, such as:

Identify the problem: Gather information from various sources, such as users, logs, or alerts, to understand the symptoms and scope of the problem.

Establish a theory of probable cause: Analyze the information and formulate one or more possible causes of the problem based on evidence or experience.

Test the theory to determine cause: Perform tests or experiments to verify or eliminate each possible cause until the root cause is found.

Establish a plan of action to resolve the problem and implement the solution: Design and execute a plan to fix the problem using appropriate tools and techniques.

Verify full system functionality and implement preventive measures: Confirm that the problem is resolved and that no other issues arise as a result of the solution. Implement preventive measures to avoid recurrence of the problem or improve performance.

Document findings, actions, and outcomes: Record the details of the problem, its cause, its solution, and its outcome for future reference or knowledge sharing. Reference: [CompTIA Server+ Certification Exam Objectives], Domain 6.0: Troubleshooting, Objective 6.1: Given a scenario involving server hardware issues (e.g., power supply failure), troubleshoot using appropriate tools.

QUESTION 87

An upper management team is investigating a security breach of the company's filesystem. It has been determined that the breach occurred within the human resources department. Which of the following was used to identify the breach in the human resources department?

- A. User groups
- B. User activity reports
- C. Password policy
- D. Multifactor authentication

Correct Answer: B

Section:

Explanation:

User activity reports were used to identify the security breach in the human resources department. User activity reports are records of the actions and events performed by users on a system or network, such as login/logout times, files accessed or modified, commands executed, or websites visited. User activity reports can help monitor and audit user behavior, detect and investigate security incidents, and enforce policies and compliance. User activity reports can be generated by various tools, such as log management software, security information and event management (SIEM) systems, or user and entity behavior analytics (UEBA) solutions. Reference: [CompTIA Server+ Certification Exam Objectives], Domain 5.0: Security, Objective 5.2: Given a scenario, apply logical access control methods.

