



CompTIA

Exam Questions SK0-005

CompTIA Server+ Certification Exam

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

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

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

Answer: B

Explanation:

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

NEW QUESTION 2

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

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

Answer: D

Explanation:

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

NEW QUESTION 3

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

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

Answer: D

Explanation:

RAID 6 is a type of RAID level that uses two parity blocks to provide fault tolerance and redundancy for data storage. RAID 6 can withstand the failure of up to two disks in the array without losing any data. RAID 6 requires a minimum of four disks to operate, and it distributes the data and parity blocks across all the disks in the array. RAID 6 has a high write penalty, which means that it takes more time and resources to write data to the disks than to read data from them. However, RAID 6 offers a high level of data protection and reliability, which makes it suitable for applications that require high availability and durability¹. RAID 1 provides redundancy and fault tolerance by mirroring the data from one disk to another disk. RAID 1 offers high read performance and data security, but it has low capacity and write performance. RAID 1 requires a minimum of two disks to operate, and it can only tolerate the failure of one disk in the array. If more than one disk fails, all the data in the array is lost². RAID 5 provides redundancy and fault tolerance by using one parity block to store information that can be used to reconstruct the data in case of a disk failure. RAID 5 requires a minimum of three disks to operate, and it distributes the data and parity blocks across all the disks in the array. RAID 5 offers a balance between performance, capacity, and data protection, but it can only tolerate the failure of one disk in the array. If more than one disk fails, all the data in the array is lost². Therefore, among these options, RAID 6 is the best solution for creating a data volume out of four disks with the most redundancy.

NEW QUESTION 4

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

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

Answer: D

Explanation:

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

NEW QUESTION 5

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

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

Answer: D

Explanation:

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

NEW QUESTION 6

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

\\server1\ITDept\

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

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

Answer: C

Explanation:

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

NEW QUESTION 7

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 8

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 9

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

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

Answer: B

Explanation:

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

NEW QUESTION 10

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

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

Answer: C

Explanation:

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

NEW QUESTION 10

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

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

Answer: B

Explanation:

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

NEW QUESTION 11

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 13

The management team has mandated the use of data-at-rest encryption for all data. Which of the following forms of encryption best achieves this goal?

- A. Drive
- B. Database
- C. Folder

D. File

Answer: A

Explanation:

Drive encryption is a form of data-at-rest encryption that encrypts the entire hard drive or solid state drive. This means that all the data on the drive, including the operating system, applications, and files, are protected from unauthorized access. Drive encryption is usually implemented at the hardware or firmware level, and requires a password, PIN, or biometric authentication to unlock the drive. Drive encryption is the most comprehensive and secure way to achieve data-at-rest encryption, as it prevents anyone from accessing the data without the proper credentials, even if they physically remove the drive from the server. References: CompTIA Server+ Study Guide, Chapter 9: Security, page 367.

NEW QUESTION 15

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

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

Answer: C

Explanation:

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

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

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

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

NEW QUESTION 16

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

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

Answer: D

Explanation:

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

NEW QUESTION 18

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

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

Answer: C

Explanation:

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

NEW QUESTION 21

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

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

Answer: C

Explanation:

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

NEW QUESTION 22

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

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

Answer: C

Explanation:

MRU, or Most Recently Used, is a concept that refers to prioritizing a connection that had previously worked successfully. It is often used in load balancing algorithms to distribute the workload among multiple servers or paths. MRU assumes that the most recently used connection is the most likely to be available and efficient, and therefore assigns the next request to that connection. This can help reduce latency and improve performance¹². The other options are incorrect because they do not refer to prioritizing a previous connection. Round robin is a concept that refers to distributing the workload equally among all available connections in a circular order¹². SCP, or Secure Copy Protocol, is a concept that refers to transferring files securely between hosts using encryption³. Link aggregation is a concept that refers to combining multiple physical links into a single logical link to increase bandwidth and redundancy⁴.

NEW QUESTION 25

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

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

Answer: D

Explanation:

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

NEW QUESTION 28

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

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

Answer: BF

Explanation:

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

NEW QUESTION 33

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

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

Answer: C

Explanation:

An emergency change request is a type of change management activity that is used to address urgent issues that pose a significant risk to the organization, such

as a system breach. An emergency change request requires immediate action and approval, and it may bypass some of the normal change management procedures, such as testing, documentation, or stakeholder communication¹².

References = 1: Change Management Plans: A Definitive Guide -Indeed(<https://www.indeed.com/career-advice/career-development/change-management-activities>) 2: The 10 Best Change Management Activities-Connecteam(<https://connecteam.com/top-10-change-management-activities/>)

NEW QUESTION 37

A systems administrator notices a newly added server cannot see any of the LUNs on the SAN. The SAN switch and the local HBA do not display any link lights. Which of the following is most likely the issue?

- A. A single-mode fiber cable is used in place of multimode.
- B. The switchport is on the wrong virtual SAN.
- C. The HBA driver needs to be installed on the server.
- D. The zoning on the fiber switch is wrong.

Answer: A

Explanation:

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

NEW QUESTION 42

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

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

Answer: C

Explanation:

An IP KVM (keyboard, video, mouse) is a device that allows remote access and control of multiple servers over a network using a web browser or a client software. An IP KVM is a cost-effective option to administer and troubleshoot network problems locally on the servers, as it eliminates the need for physical presence or dedicated hardware for each server. A management port (A) is a network interface that is used for out-of-band management of network devices, such as routers or switches. A management port does not provide local access to servers. A crash cart (B) is a mobile unit that contains a monitor, keyboard, mouse, and other tools for troubleshooting servers in a data center. A crash cart requires physical access to each server and may not be cost-effective for many racks of servers. A KVM (D) is a device that allows switching between multiple servers using a single keyboard, video, and mouse. A KVM does not provide remote access over a network and requires physical connection to each server. References: <https://www.enterprisestorageforum.com/management/best-data-storage-solutions-and-software-2021/><https://www.microsoft.com/en-us/microsoft-365/business-insights-ideas/resources/cloud-storage-vs-on-premises-servers>

NEW QUESTION 45

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

- A. Unplug the power cables from the redundant power supplies, leaving just the minimum required.
- B. Convert the physical servers to the hypervisors and retire the ten servers.
- C. Reimage the physical servers and retire all ten servers after the migration is complete.
- D. Convert the ten servers to power-efficient core editions.

Answer: B

Explanation:

This option would reduce power consumption and maintenance costs by consolidating the physical servers into virtual machines on the hypervisors. This would also free up space and resources in the data center. The other options would either not address the request, increase power consumption, or require more maintenance.

NEW QUESTION 48

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

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

Answer: A

Explanation:

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

NEW QUESTION 53

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

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

Answer: B

Explanation:

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

NEW QUESTION 55

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

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

Answer: C

Explanation:

RAID (Redundant Array of Independent Disks) is a technology that combines multiple hard drives into a logical unit that provides improved performance, reliability, or capacity. RAID can be implemented by hardware, software, or a combination of both. Hardware RAID uses a dedicated controller to manage the RAID array, while software RAID uses the operating system or a driver to do the same¹.

In this scenario, the technician noted that the RAID hard drives were functional while troubleshooting a motherboard failure. This means that the data on the drives was not corrupted or lost. However, the technician installed a spare motherboard with similar specifications and used the original components. This means that the new motherboard may not have the same RAID configuration as the old one, or it may not recognize the existing RAID array at all. Therefore, the technician needs to reconfigure the RAID in order to restore operations with minimal downtime.

NEW QUESTION 57

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

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

Answer: B

Explanation:

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

NEW QUESTION 58

An administrator has been asked to disable CPU hyperthreading on a server to satisfy a licensing issue. Which of the following best describes how the administrator will likely perform this action?

- A. Use a RDP/VNC session.
- B. Modify the startup configuration.
- C. Use a PowerShell/Bash script.
- D. Use the BIOS/UEFI setup.

Answer: D

Explanation:

The BIOS (Basic Input/Output System) or UEFI (Unified Extensible Firmware Interface) setup is a program that allows users to configure the hardware settings of a computer, such as the CPU, memory, disk, and boot options. The BIOS/UEFI setup can be accessed by pressing a specific key (such as F2, F10, or Delete) during the boot process, before the operating system loads¹².

One of the settings that can be changed in the BIOS/UEFI setup is the CPU hyperthreading option. Hyperthreading is a technology that enables a single physical CPU core to execute two threads or tasks simultaneously, improving the performance and efficiency of multi-threaded applications. However, some software licenses may limit the number of CPU cores or threads that can be used, and therefore require disabling hyperthreading on the server³⁴.

To disable hyperthreading on a server, the administrator will likely need to enter the BIOS/UEFI setup and navigate to the processor options menu. There, the administrator will find a setting for Intel® Hyperthreading Technology or Hyperthreading Function, which can be enabled or disabled. The administrator will need to disable this setting and save the changes. This will turn off hyperthreading on the server and reduce the number of logical CPUs to match the number of physical cores⁵.

NEW QUESTION 61

An application server cannot communicate with a newly installed database server. The database server, which has static IP information, is reading the following output from ipconf ig:

```
IP: 10.0.10.240
Mask: 255.255.255.128
Gateway: 10.0.10.1
```

The application server is reading the following output from ipconf ig:

```
IP: 10.0.10.25
Mask: 255.255.255.128
Gateway: 10.0.10.1
```

Which of the following most likely contains an error?

- A. IP address
- B. DHCP
- C. Gateway
- D. Subnet mask

Answer: D

Explanation:

The subnet mask is most likely containing an error that prevents the application server from communicating with the newly installed database server. The subnet mask is a binary number that defines how many bits of an IP address are used for the network portion and how many bits are used for the host portion. The subnet mask determines which devices belong to the same network or subnet and can communicate directly with each other without routing or switching devices. The subnet mask of the database server is 255.255.O.O, which means that all 32 bits of its IP address are used for the network portion and none for the host portion, which is invalid and makes it unreachable by any other device on any network or subnet. The subnet mask of the application server is 255.O.O.O, which means that only 8 bits of its IP address are used for the network portion and 24 bits are used for the host portion, which is also uncommon and makes it incompatible with most networks or subnets. References: [CompTIA Server+ Certification Exam Objectives], Domain 4.0: Networking, Objective 4.1: Given a scenario, configure network settings for servers.

NEW QUESTION 64

An administrator is deploying a new secure web server. The only administration method that is permitted is to connect via RDP. Which of the following ports should be allowed? (Select TWO).

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

Answer: DF

Explanation:

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

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

NEW QUESTION 67

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

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

Answer: BE

Explanation:

These are two measures that should be taken when a data breach occurs. A data breach is an unauthorized or illegal access to confidential or sensitive data by an internal or external actor. A data breach can result in financial losses, reputational damage, legal liabilities, and regulatory penalties for the affected organization. Disclosing the incident is a measure that involves informing the relevant stakeholders, such as customers, employees, partners, regulators, and law enforcement, about the nature, scope, and impact of the data breach. Disclosing the incident can help to mitigate the negative consequences of the data breach, comply with legal obligations, and restore trust and confidence. Identifying the exploited vulnerability is a measure that involves investigating and analyzing the root

cause and source of the data breach. Identifying the exploited vulnerability can help to prevent further data loss, remediate the security gaps, and improve the security posture of the organization. Restoring the data from backup is a measure that involves recovering the lost or corrupted data from a secondary storage device or location. However, this does not address the underlying issue of how the data breach occurred or prevent future breaches. Disabling unnecessary ports is a measure that involves closing or blocking network communication endpoints that are not required for legitimate purposes. However, this does not address how the data breach occurred or what vulnerability was exploited. Running an antivirus scan is a measure that involves detecting and removing malicious software from a system or network. However, this does not address how the data breach occurred or what vulnerability was exploited. Moving the data to a different location is a measure that involves transferring the data to another storage device or location that may be more secure or less accessible. However, this does not address how the data breach occurred or what vulnerability was exploited. References: <https://www.howtogeek.com/428483/what-is-end-to-end-encryption-and-why-does-it-matter/> <https://www.howtogeek.com/202794/what-is-the-difference-between-127.0.0.1-and-0.0.0.0/> <https://www.howtogeek.com/443611/how-to-encrypt-your-macs-system-drive-removable-devices-and-individual-files/>

NEW QUESTION 69

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

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

Answer: B

Explanation:

The port that must remain open for a secure website is port 443. Port 443 is used by Hypertext Transfer Protocol Secure (HTTPS), which is an extension of HTTP that encrypts and authenticates the communication between a web server and a web browser. HTTPS ensures that the data transmitted over the web is protected from eavesdropping, tampering, or spoofing. Therefore, port 443 must remain open for a secure website to function properly.

References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 2, Lesson 2.2, Objective 2.2

NEW QUESTION 72

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

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

Answer: D

Explanation:

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

NEW QUESTION 75

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

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

Answer: C

Explanation:

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

Reference:

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

NEW QUESTION 77

An administrator is tasked with building an environment consisting of four servers that can each serve the same website. Which of the following concepts is described?

- A. Load balancing
- B. Direct access
- C. Overprovisioning
- D. Network teaming

Answer: A

Explanation:

Load balancing is a concept that distributes the workload across multiple servers or other resources to optimize performance, availability, and scalability. Load balancing can be implemented at different layers of the network, such as the application layer, the transport layer, or the network layer. Load balancing can use various algorithms or methods to determine how to distribute the traffic, such as round robin, least connections, or weighted distribution.

References: CompTIA Server+ Study Guide, Chapter 6: Networking, page 241.

NEW QUESTION 79

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

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

Answer: B

Explanation:

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

NEW QUESTION 84

A security technician generated a public/private key pair on a server. The technician needs to copy the key pair to another server on a different subnet. Which of the following is the most secure method to copy the keys?

? HTTP

- A. FTP
- B. SCP
- C. USB

Answer: C

Explanation:

SCP (Secure Copy Protocol) is a protocol that allows users to securely transfer files between servers using SSH (Secure Shell) encryption. SCP encrypts both the data and the authentication information, preventing unauthorized access, interception, or modification of the files. SCP also preserves the file attributes, such as permissions, timestamps, and ownership.

NEW QUESTION 89

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

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

Answer: A

Explanation:

To make a new drive available on a server, the administrator needs to partition the drive first. Partitioning is a process that divides the drive into one or more logical sections that can be formatted and assigned drive letters or mount points. Partitioning can be done using tools such as Disk Management on Windows or fdisk on Linux. Creating a new disk quota would not help, as disk quotas are used to limit the amount of disk space that users or groups can use on a partition. Configuring the drive as dynamic would not help either, as dynamic disks are used to create volumes that span multiple disks or use RAID features. Setting the compression would not help, as compression is used to reduce the size of files on a partition. References: <https://www.howtogeek.com/school/using-windows-admin-tools-like-a-pro/lesson2/> <https://www.howtogeek.com/howto/17001/how-to-format-a-usb-drive-in-ubuntu-using-gparted/>

NEW QUESTION 94

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

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

Answer: C

Explanation:

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

NEW QUESTION 96

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

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

Answer: C

Explanation:

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

NEW QUESTION 101

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

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

Answer: B

Explanation:

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

NEW QUESTION 103

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

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

Answer: A

Explanation:

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

NEW QUESTION 105

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

- A. Network encapsulation
- B. Off-site data
- C. Secure FTP
- D. Data in transit

Answer: D

Explanation:

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

NEW QUESTION 109

Which of the following license types most commonly describes a product that incurs a yearly cost regardless of how much it is used?

- A. Physical
- B. Subscription
- C. Open-source
- D. Per instance
- E. Per concurrent user

Answer: B

Explanation:

A subscription license is a type of license that grants the user the right to use a product or service for a fixed period of time, usually a year. The user pays a recurring fee, regardless of how much they use the product or service. Subscription licenses are common for cloud-based software and services, such as Microsoft 365 or DocuSign2.

References = 1: Compare All Microsoft 365 Plans (Formerly Office 365) - Microsoft Store(<https://www.microsoft.com/en-us/microsoft-365/buy/compare-all-microsoft-365-products>) 2: DocuSign Pricing | eSignature Plans for Personal & Business(<https://ecom.docusign.com/plans-and-pricing/esignature>)

NEW QUESTION 114

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

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

Answer: D

Explanation:

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

NEW QUESTION 116

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

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

Answer: C

Explanation:

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

NEW QUESTION 117

A technician recently replaced a NIC that was not functioning. Since then, no device driver is found when starting the server, and the network card is not functioning. Which of the following should the technician check first?

- A. The boot log
- B. The BIOS
- C. The HCL
- D. The event log

Answer: C

Explanation:

The technician should check the hardware compatibility list (HCL) first to see if the new NIC is supported by the server's operating system. The HCL is a list of hardware devices that have been tested and verified to work with a specific operating system. If the NIC is not on the HCL, it means that there is no device driver available or compatible for it, and the NIC will not function properly.

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

NEW QUESTION 122

A technician is tasked with upgrading 24 hosts simultaneously with a Type 1 hypervisor. Which of the following protocols should the technician use for this upgrade?

- A. VPN

- B. TFTP
- C. SSH
- D. HTTP

Answer: B

Explanation:

TFTP (Trivial File Transfer Protocol) is a simple and lightweight protocol that can be used to transfer files over a network. TFTP is often used to upgrade firmware or software on network devices, such as routers, switches, or servers. TFTP can also be used to install a Type 1 hypervisor, such as VMware ESXi, on multiple hosts simultaneously¹². References = 1: How to Install VMware ESXi Type 1 Hypervisor - MatthewEaton.net(<https://mattheweaton.net/posts/how-to-install-vmware-esxi-type-1-hypervisor/>) 2: Explore Type 1 Hypervisors - Set Up Virtual Machines Using VirtualBox and vSphere - OpenClassrooms(<https://openclassrooms.com/en/courses/7163136-set-up-virtual-machines-using-virtualbox-and-vsphere/7358546-explore-type-1-hypervisors>)

NEW QUESTION 123

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

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

Answer: D

Explanation:

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

NEW QUESTION 126

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

- * 1. The OS must be installed in a separate disk partition. In case of hard drive failure, it cannot be affected.
- * 2. Application data IOPS performance is a must.
- * 3. Data availability is a high priority, even in the case of multiple hard drive failures.

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

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

Answer: BDE

Explanation:

To comply with the user requirements, the best options are to install the OS on a RAID 1 array, configure RAID 5 for the application data, and use SSD hard drives for the data application array. Here is why:

? RAID 1 is a mirroring technique that creates an exact copy of data on two disks.

This provides redundancy and fault tolerance in case of hard drive failure. RAID 1 also improves read performance since either disk can be read at the same time. Therefore, installing the OS on a RAID 1 array meets the first requirement of separating the OS from the application data and protecting it from hard drive failure.

? RAID 5 is a striping technique with parity that distributes data and parity blocks

across three or more disks. This provides improved performance and storage efficiency compared to RAID 1, as well as fault tolerance in case of a single disk failure. Therefore, configuring RAID 5 for the application data meets the second and third requirements of providing high IOPS performance and data availability.

? SSD hard drives are solid-state drives that use flash memory to store data. They

have no moving parts and offer faster read and write speeds, lower latency, and lower power consumption than traditional HDDs. Therefore, using SSD hard drives for the data application array meets the second requirement of providing high IOPS performance.

References:

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

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

NEW QUESTION 128

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

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

Answer: D

Explanation:

Maintenance and support is a type of license that would most likely include vendor assistance. Maintenance and support is a contract that defines the level and scope of service and assistance that a vendor provides to a customer for using their software product. Maintenance and support may include technical support, bug fixes, patches, updates, upgrades, documentation, training, and other benefits. Maintenance and support licenses usually have an annual fee based on the

number of users or devices covered by the contract. Open-source is a type of license that allows free access to the source code and modification and distribution of the software product, but does not guarantee vendor assistance. Version compatibility is not a type of license, but a feature that ensures software products can work with different versions of operating systems or other software products. Subscription is a type of license that allows access to software products for a limited period of time based on recurring payments, but does not necessarily include vendor assistance. References: <https://www.techopedia.com/definition/1440/software-licensing>
<https://www.techopedia.com/definition/1032/business-impact-analysis-bia>

NEW QUESTION 130

A systems administrator is setting up a server farm for a new company. The company has a public range of IP addresses and uses the addresses internally. Which of the following IP addresses best fits this scenario?

- A. 10.3.7.27
- B. 127.0.0.1
- C. 192.168.7.1
- D. 216.176.128.10

Answer: D

Explanation:

The IP address that best fits this scenario is 216.176.128.10. This is a public IP address that belongs to a range of addresses that are assigned and registered by an Internet service provider (ISP) and can be accessed from anywhere on the Internet. The company has a public range of IP addresses and uses them internally, which means that they do not use private IP addresses or network address translation (NAT) to communicate within their network.

References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 2, Lesson 2.2, Objective 2.2

NEW QUESTION 132

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

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

Answer: B

Explanation:

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

NEW QUESTION 136

A data center environment currently hosts more than 100 servers that include homegrown and commercial software. The management team has asked the server administrator to find a way to eliminate all company-owned data centers. Which of the following models will the administrator most likely choose to meet this need?

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

Answer: C

Explanation:

A public cloud model will most likely meet the need of eliminating all company-owned data centers. A public cloud is a type of cloud computing service that is provided by a third-party vendor over the internet. A public cloud offers scalability, flexibility, and cost-effectiveness for hosting servers and applications, as the customers only pay for the resources they use and do not have to maintain their own infrastructure. A public cloud can also provide high availability, security, and performance for the servers and applications, as the vendor manages the underlying hardware and software. A public cloud can support various types of services, such as software as a service (SaaS), platform as a service (PaaS), or infrastructure as a service (IaaS). References: [CompTIA Server+ Certification Exam Objectives], Domain 1.0: Server Administration, Objective 1.2: Given a scenario, compare and contrast server roles and requirements for each.

NEW QUESTION 137

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

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

Answer: BE

Explanation:

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

References:

CompTIA Server+ Certification Exam Objectives1, page 12

NEW QUESTION 140

Which of the following physical security concepts would most likely be used to limit personnel access to a restricted area within a data center?

- A. An access control vestibule
- B. Video surveillance
- C. Bollards
- D. Data center camouflage

Answer: A

Explanation:

An access control vestibule is a physical security concept that limits personnel access to a restricted area within a data center. It is a small room or hallway that has two doors: one that leads to the outside and one that leads to the restricted area. The doors are controlled by an electronic lock that requires authentication, such as a card reader, biometric scanner, or keypad. Only authorized personnel can enter the vestibule and access the restricted area. References: CompTIA Server+ Certification Exam Objectives, Domain 5.0: Security, Objective 5.1: Given a scenario, apply physical security methods to a server.

NEW QUESTION 141

An analyst is planning a new point-to-point Ethernet connection between a university campus and a newly acquired space downtown that is about 5mi (8km) away. The analyst would like the fastest possible connection speed. Which of the following would best meet the analyst's needs?

- A. 1000BASE-LX 1Gb single-mode plenum fiber connection
- B. 10GBASE-T 10Gb copper plenum Ethernet connection
- C. 1000BASE-T 1Gb copper non-plenum Ethernet connection
- D. 10GBASE-SR 10Gb multimode plenum fiber connection

Answer: A

Explanation:

A 1000BASE-LX 1Gb single-mode plenum fiber connection would best meet the analyst's needs for a new point-to-point Ethernet connection between a university campus and a newly acquired space downtown that is about 5mi (8km) away. A 1000BASE-LX is a type of Ethernet standard that supports data transmission at 1 gigabit per second over single-mode fiber cables using long wavelength lasers. A single-mode fiber cable is a type of optical fiber cable that has a small core diameter and allows only one mode of light to propagate through it. A single-mode fiber cable can transmit data over long distances at high speeds, but it requires more expensive transceivers and connectors than multimode fiber cables. A plenum fiber cable is a type of optical fiber cable that has a special coating that prevents the spread of fire or toxic fumes in case of burning. A plenum fiber cable is suitable for installation in plenum spaces, which are areas used for air circulation in buildings, such as above ceilings or below floors. References: [CompTIA Server+ Certification Exam Objectives], Domain 4.0: Networking, Objective 4.2: Given a scenario involving server networking issues (e.g., network interface card failure), troubleshoot using appropriate tools.

NEW QUESTION 143

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

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

Answer: C

Explanation:

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

NEW QUESTION 147

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

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

Answer: A

Explanation:

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

NEW QUESTION 148

A systems administrator is trying to determine why users in the human resources department cannot access an application server. The systems administrator reviews the application logs but does not see any attempts by the users to access the application. Which of the following is preventing the users from accessing the application server?

- A. NAT
- B. ICMP

- C. VLAN
- D. NIDS

Answer: C

Explanation:

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

NEW QUESTION 151

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

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

Answer: CE

Explanation:

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

NEW QUESTION 153

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

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

Answer: B

Explanation:

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

NEW QUESTION 157

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

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

Answer: A

Explanation:

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

NEW QUESTION 158

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

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

Answer: A

Explanation:

The most likely cause of the issue is that the disk uses GPT. GPT stands for GUID Partition Table, which is a newer standard for disk partitioning that supports larger disks and more partitions than the older MBR (Master Boot Record) standard. However, GPT is not compatible with some older operating systems, such as Windows XP or Windows Server 2003. Therefore, if the data drive was formatted with GPT on a new Windows server and then moved to an older Windows server, the older server may not be able to recognize the GPT partitions and access the data on the drive.

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

NEW QUESTION 163

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

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

Answer: D

Explanation:

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

NEW QUESTION 167

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

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

Answer: C

Explanation:

RTO (Recovery Time Objective) is a 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. References: <https://parachute.cloud/rto-vs-rpo/> <https://www.techopedia.com/definition/13622/service-level-agreement-sla> <https://www.techopedia.com/definition/1032/business-impact-analysis-bia> <https://www.techopedia.com/definition/8239/mean-time-to-repair-mttr>

NEW QUESTION 171

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

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

Answer: D

Explanation:

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

NEW QUESTION 172

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

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

Answer: AE

Explanation:

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

NEW QUESTION 174

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

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

Answer: D

Explanation:

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

NEW QUESTION 178

A security analyst completed a port scan of the corporate production-server network. Results of the scan were then provided to a systems administrator for immediate action. The following table represents the requested changes:

Server name	Block	Do not change
MailSrv	20, 21, 22, 23, 53	25, 3389
WebSrv	20, 21, 22, 23, 53	80, 443, 3389
SQLSrv	20, 21, 22, 23, 53	1443, 3389
DNSSrv	20, 21, 22, 23, 53	67, 68, 3389

The systems administrator created local firewall rules to block the ports indicated above. Immediately, the service desk began receiving calls about the internet being down. The systems administrator then reversed the changes, and the internet became available again. Which of the following ports on DNSSrv must remain open when the firewall rules are reapplied?

- A. 20
- B. 21
- C. 22
- D. 23
- E. 53

Answer: E

Explanation:

Port 53 is the standard port for DNS (Domain Name System) queries and responses. DNS is a service that translates domain names (such as www.example.com) into IP addresses (such as 192.0.2.1) and vice versa. DNS is essential for internet connectivity, as it allows users and applications to access websites and other online resources by using human-readable names instead of numerical addresses.

The DNSSrv server is a DNS server that provides name resolution for the corporate network. If port 53 is blocked on this server, it will not be able to communicate with other DNS servers or clients, and the name resolution will fail. This will prevent users from accessing any websites or online services that rely on domain names, such as web browsers, email clients, or cloud applications. Therefore, port 53 must remain open on DNSSrv to allow DNS traffic to flow.

NEW QUESTION 181

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

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

Answer: A

Explanation:

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

NEW QUESTION 184

A site is considered a warm site when it:
 ? has basic technical facilities connected to it.
 ? has faulty air conditioning that is awaiting service.
 ? is almost ready to take over all operations from the primary site.

- A. is fully operational and continuously providing services.

Answer: A

Explanation:

A warm site is a backup site that has some of the necessary hardware, software, and network resources to resume operations, but not all of them. A warm site requires some time and effort to become fully operational. A warm site is different from a cold site, which has minimal or no resources, and a hot site, which has all the resources and is ready to take over immediately.

References: CompTIA Server+ Study Guide, Chapter 10: Disaster Recovery, page 403.

NEW QUESTION 188

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

777

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

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

Answer: A

Explanation:

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

NEW QUESTION 192

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

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

Answer: B

Explanation:

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

NEW QUESTION 196

A technician is sizing a new server and, for service reasons, needs as many hot-swappable components as possible. Which of the following server components can most commonly be replaced without downtime? (Select three).

- A. Drives
- B. Fans
- C. CMOSIC
- D. Processor
- E. Power supplies
- F. Motherboard
- G. Memory
- H. BIOS

Answer: ABE

Explanation:

Drives, fans, and power supplies are server components that can most commonly be replaced without downtime if they are hot-swappable. Hot-swappable components can be removed and inserted while the server is running, without affecting its operation or performance. Drives store data and applications, fans cool down the server components, and power supplies provide electricity to the server. Replacing these components can prevent data loss, overheating, or power failure. References: CompTIA Server+ Certification Exam Objectives, Domain 2.0: Hardware, Objective 2.2: Given a scenario, install, configure and maintain server components.

NEW QUESTION 197

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

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

Answer: A

Explanation:

The administrator should use an IP KVM to access the server's console remotely for bare-metal maintenance. An IP KVM stands for Internet Protocol Keyboard Video Mouse, which is a device that allows remote control of a server's keyboard, video, and mouse over a network connection, such as LAN or Internet. An IP KVM enables an administrator to perform tasks such as BIOS configuration, boot sequence selection, operating system installation, etc., without being physically present at the server location. The other options are not suitable for bare-metal maintenance because they require either physical access to the server (a crash cart) or an operating system running on the server (VNC, RDP, SSH). A crash cart is a mobile unit that contains a monitor, keyboard, mouse, and cables that can

be plugged into a server for direct access to its console. VNC stands for Virtual Network Computing, which is a software that allows remote desktop sharing and control over a network connection using a graphical user interface (GUI). RDP stands for Remote Desktop Protocol, which is a protocol that allows remote desktop access and control over a network connection using a GUI or command-line interface (CLI). SSH stands for Secure Shell, which is a protocol that allows secure remote login and command execution over a network connection using a CLI.

NEW QUESTION 199

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

Answer: B

Explanation:

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

NEW QUESTION 201

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

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

Answer: A

Explanation:

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

NEW QUESTION 205

A user can successfully connect to a database server from a home office but is unable to access it from a hotel room. Which of the following authentication methods is most likely configured?

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

Answer: D

Explanation:

Scope-based authentication is a method of restricting access to resources based on the location, network, or device of the user. It can be used to prevent unauthorized access from outside the organization's network or from untrusted devices. In this case, the user can connect to the database server from the home office, which is likely within the scope of the authentication policy, but not from the hotel room, which is outside the scope.

References:

CompTIA Server+ Certification Exam Objectives1, page 15 CompTIA Server+: Authentication & Authorization2

NEW QUESTION 207

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

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

Answer: B

Explanation:

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

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

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

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

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

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

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

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

NEW QUESTION 208

HOTSPOT


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

INSTRUCTIONS

Perform the following steps:

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

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



Proxy Server Routing Table			
Destination	Netmask	Gateway	Interface
0.0.0.0	0.0.0.0	▼	▼
		192.168.3.0	192.168.4.1
		192.168.4.0	192.168.1.1
		192.168.1.1	192.168.3.0
		192.168.2.0	192.168.1.0
		192.168.1.0	192.168.2.2
		192.168.4.1	0.0.0.0
		192.168.2.1	192.168.3.1
		0.0.0.0	255.255.255.0
		192.168.3.1	192.168.3.2
		255.255.255.0	192.168.4.0
		192.168.3.2	192.168.2.1
		192.168.2.2	192.168.2.0
		192.168.1.0	255.255.255.0
192.168.3.0	192.168.4.1		
192.168.4.0	192.168.1.1		
192.168.1.1	192.168.3.0		
192.168.2.0	192.168.1.0		
192.168.1.0	192.168.2.2		
192.168.4.1	0.0.0.0		
192.168.2.1	192.168.3.1		
0.0.0.0	255.255.255.0		
192.168.3.1	192.168.3.2		
255.255.255.0	192.168.4.0		
192.168.3.2	192.168.2.1		
192.168.2.2	192.168.2.0		

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Proxy Server Routing Table			
Destination	Netmask	Gateway	Interface
0.0.0.0	0.0.0.0	▼	▼
		192.168.3.0	192.168.4.1
		192.168.4.0	192.168.1.1
		192.168.1.1	192.168.3.0
		192.168.2.0	192.168.1.0
		192.168.1.0	192.168.2.2
		192.168.4.1	0.0.0.0
		192.168.2.1	192.168.3.1
		0.0.0.0	255.255.255.0
		192.168.3.1	192.168.3.2
		255.255.255.0	192.168.4.0
		192.168.3.2	192.168.2.1
		192.168.2.2	192.168.2.0
		192.168.1.0	255.255.255.0
192.168.3.0	192.168.4.1		
192.168.4.0	192.168.1.1		
192.168.1.1	192.168.3.0		
192.168.2.0	192.168.1.0		
192.168.1.0	192.168.2.2		
192.168.4.1	0.0.0.0		
192.168.2.1	192.168.3.1		
0.0.0.0	255.255.255.0		
192.168.3.1	192.168.3.2		
255.255.255.0	192.168.4.0		
192.168.3.2	192.168.2.1		
192.168.2.2	192.168.2.0		

NEW QUESTION 213

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

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

Answer: B

Explanation:

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

NEW QUESTION 214

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

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

Answer: B

Explanation:

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

References:

? CompTIA Server+ Certification Exam Objectives1, page 14

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

NEW QUESTION 218

A remote, embedded IoT server is having a Linux OS upgrade installed. Which of the following is the best method to stage the new media for the default boot device of the server?

- A. Copy and send an SSD to the site.
- B. Copy and send a DVD to the site.
- C. Copy and send a SATA drive to the site.
- D. Copy and send a microSD card to the site.

Answer: D

Explanation:

A microSD card is the best method to stage the new media for the default boot device of a remote embedded IoT server that is having a Linux OS upgrade installed. A microSD card is a small and portable storage device that can store large amounts of data. It can be easily inserted into the slot of an embedded IoT server, which is a small and low-power device that performs specific tasks and connects to other devices over a network. A microSD card can also be formatted with different file systems, such as FAT32 or ext4, which are compatible with Linux OS. References: CompTIA Server+ Certification Exam Objectives, Domain 4.0: Networking, Objective 4.3: Given a scenario, configure servers for IoT applications.

NEW QUESTION 219

A company needs a media server set up that provides the highest availability with a minimum requirement of at least 10TB. The company purchased five HDDs, each with a 4TB capacity. Which of the options would provide the highest fault tolerance and meet the requirements?

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

Answer: C

Explanation:

RAID 6 is a RAID level that uses disk striping with two parity blocks distributed across all member disks. It can tolerate the failure of up to two disks without losing any data. RAID 6 can provide a minimum of 10TB of usable storage space with five 4TB disks, as the formula for calculating the RAID 6 capacity is $(n-2) \times S_{min}$, where n is the number of disks and S_{min} is the smallest disk size. In this case, the RAID 6 capacity is $(5-2) \times 4TB = 12TB$. References:

? CompTIA Server+ Certification Exam Objectives1, page 8

? RAID Levels and Types Explained: Advantages and Disadvantages2

? RAID Levels & Fault Tolerance3

NEW QUESTION 223

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

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

Answer: B

Explanation:

The output of the configuration shows that RAID 5 was configured properly using four 200GB drives. The approximate volume size of 400GB is correct, since RAID 5 uses one disk for parity and the rest for data. Therefore, the usable storage capacity is three-fourths of the total capacity, which is 600GB out of 800GB.

The other RAID levels given would result in different volume sizes: RAID 0 would result in 800GB, RAID 1 would result in 200GB, and JBOD would result in an error since it does not support multiple drives in a single volume. References: https://en.wikipedia.org/wiki/Standard_RAID_levels#RAID_5

NEW QUESTION 226

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

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

Answer: DE

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.

NEW QUESTION 231

Which of the following BEST describes a warm site?

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

Answer: D

Explanation:

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

NEW QUESTION 233

Which of the following security risks provides unauthorized access to an application?

- A. Backdoor
- B. Data corruption
- C. Insider threat
- D. Social engineering

Answer: A

Explanation:

A backdoor is a security risk that provides unauthorized access to an application. A backdoor is a hidden or undocumented way of bypassing the normal authentication or encryption mechanisms of an application, allowing an attacker to gain remote access, execute commands, or steal data. A backdoor can be created intentionally by the developer, maliciously by an attacker, or unintentionally by a programming error. References: CompTIA Server+ Certification Exam Objectives, Domain 5.0: Security, Objective 5.2: Given a scenario, apply logical access control methods.

NEW QUESTION 235

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

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

Answer: B

Explanation:

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

NEW QUESTION 239

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

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

Answer: BD

Explanation:

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

NEW QUESTION 244

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

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

Answer: A

Explanation:

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

NEW QUESTION 247

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 248

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

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

Answer: D

Explanation:

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

NEW QUESTION 252

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

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

Answer: C

Explanation:

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

NEW QUESTION 253

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

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

Answer: B

Explanation:

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

NEW QUESTION 255

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

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

Answer: A

Explanation:

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

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

Reference: <https://www.ibm.com/cloud/learn/recovery-time-objective>

NEW QUESTION 256

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

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

Answer: C

Explanation:

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

NEW QUESTION 260

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 261

A user has been unable to authenticate to the company's external, web-based database after clicking a link in an email that required the user to change the account password. Which of the following steps should the company take next?

- A. Disable the user's account and inform the security team.
- B. Create a new log-in to the external database.
- C. Ask the user to use the link again to reset the password.
- D. Reset the user's password and ask the user to log in again.

Answer: A

Explanation:

The user has likely fallen victim to a phishing scam, which is a fraudulent attempt to obtain sensitive information, such as passwords, by disguising as a legitimate entity. The link in the email that required the user to change the account password was probably a fake website that mimicked the company's external database, and captured the user's credentials when they entered them. This could compromise the security and integrity of the company's data, as well as the user's identity and privacy¹².

The company should take immediate action to prevent further damage and investigate the incident. The first step is to disable the user's account and inform the security team. Disabling the user's account can prevent unauthorized access to the external database by the attackers, who may use the stolen credentials to log in and manipulate or steal data. Informing the security team can alert them of the breach and allow them to take appropriate measures, such as scanning for malware, changing passwords, notifying other users, and reporting the incident³⁴.

NEW QUESTION 266

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

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

Answer: CF

Explanation:

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

NEW QUESTION 267

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 271

.....

Relate Links

100% Pass Your SK0-005 Exam with Exam Bible Prep Materials

<https://www.exambible.com/SK0-005-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>