GetValidTest

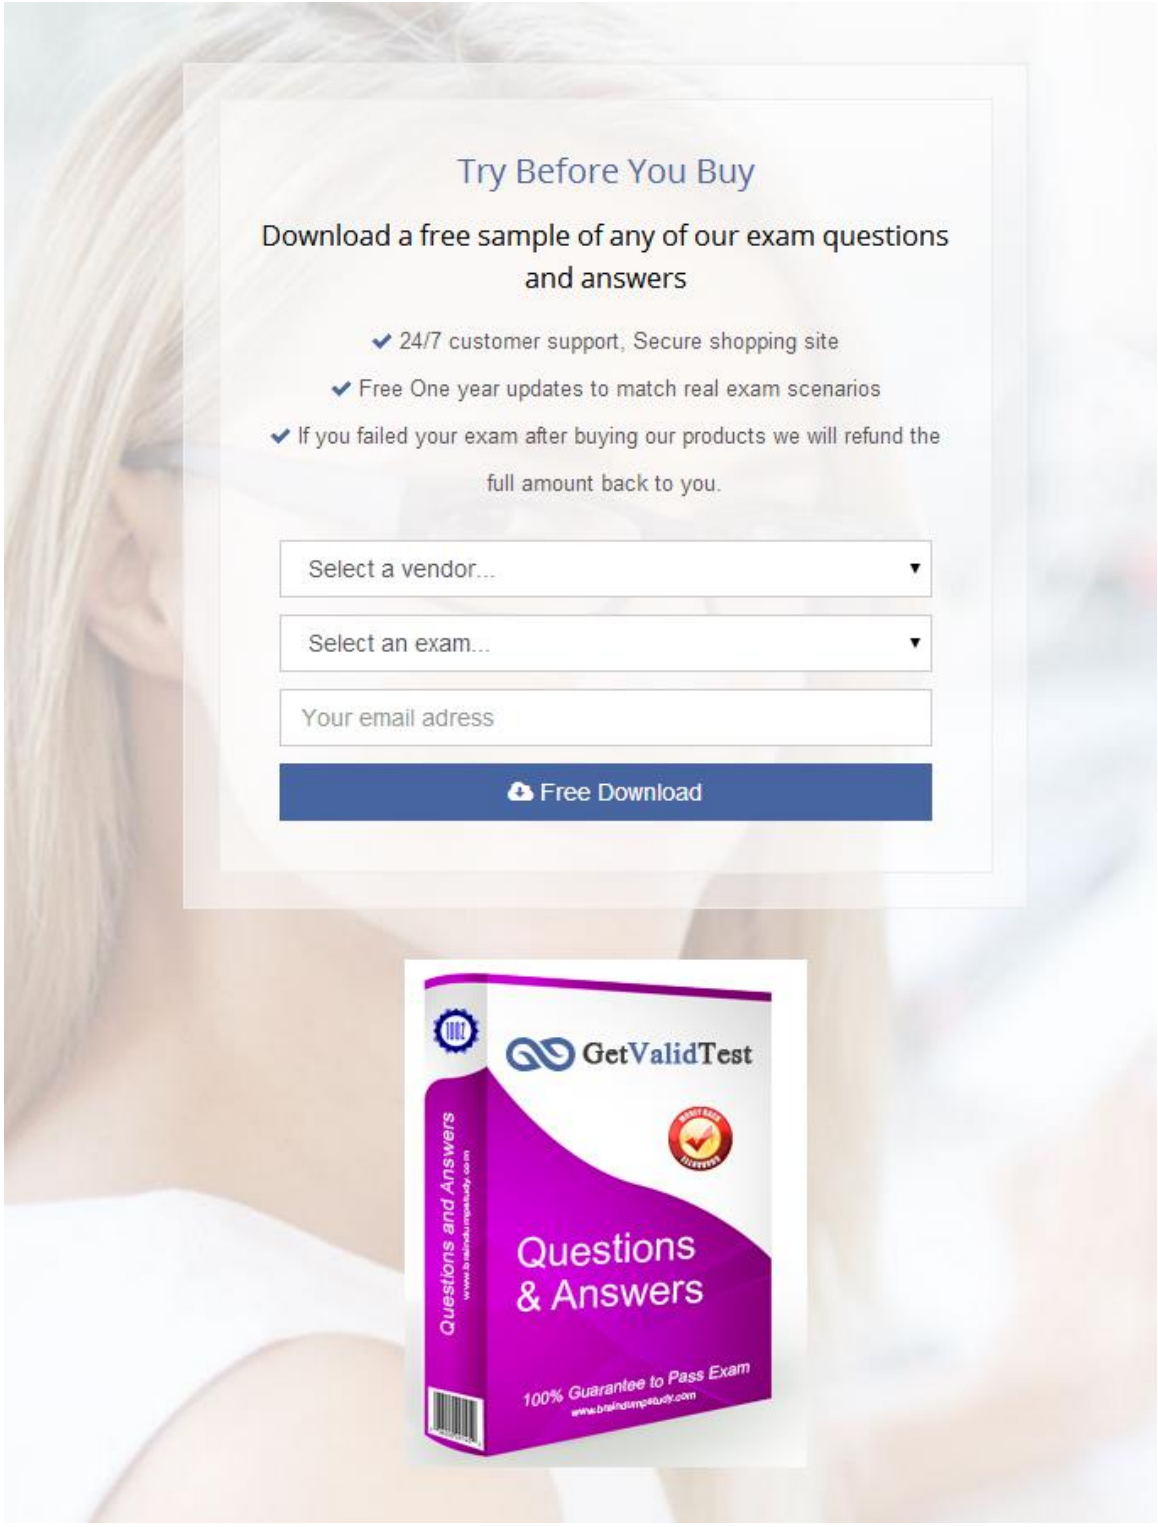

http://www.getvalidtest.com

GetValidTest - Valid test study guide, get certification

Exam : **SK0-005**

Title : CompTIA Server+ Certification Exam

Vendor : CompTIA

Version : DEMO

NO.1 A server administrator is installing a new server that uses 40G0 network connectivity. The administrator needs to find the proper cables to connect the server to the switch. Which of the following connectors should the administrator use?

A. SFP+

B. GBIC

C. SFP

D. QSFP+

Answer: D

Explanation:

QSFP+ is a type of connector that should be used to connect a server to a switch that uses 40G network connectivity.

QSFP+ (Quad Small Form-factor Pluggable Plus) is a compact, hot-pluggable transceiver module that supports data rates up to 40 Gbps.

QSFP+ modules can be used for various network protocols and media types, such as Ethernet, Fibre Channel, InfiniBand, or optical fiber.

QSFP+ modules have a 38-pin edge connector and can be inserted into a QSFP+ port on a switch or a server. SFP+ (Small Form-factor Pluggable Plus) is a type of connector that supports data rates up to 10 Gbps, but not 40 Gbps. SFP+ modules have a 20-pin edge connector and can be inserted into an SFP+ port on a switch or a server. GBIC (Gigabit Interface Converter) is an older type of connector that supports data rates up to 1 Gbps, but not 40 Gbps.

GBIC modules have an SC duplex connector and can be inserted into a GBIC port on a switch or a server.

SFP (Small Form-factor Pluggable) is another older type of connector that supports data rates up to 1 Gbps or 4 Gbps, but not 40 Gbps.

SFP modules have an LC duplex connector and can be inserted into an SFP port on a switch or a server.

References: https://www.howtogeek.com/190014/virtualization-basics-understanding-techniques-and-fundamentals/

https://www.howtogeek.com/428483/what-is-end-to-end-encryption-and-why- does-it-matter/

https://www.howtogeek.com/443611/how-to-encrypt-your-macs-system-drive-removable-devices-and-individual-files/

NO.2 A senior administrator instructs a technician to run the following script on a Linux server:

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

The script mostly returns the following message: Connection refused.

However, there are several entries in the console display that look like this:

80

Connected to localhost

443

Connected to localhost

Which of the following actions should the technician perform NEXT?

A. Look for an unauthorized HTTP service on this server

B. Look for a virus infection on this server

C. Look for an unauthorized Telnet service on this server

D. Look for an unauthorized port scanning service on this server.

Answer: A
Explanation:
The script that the technician is running is trying to connect to every port on the localhost (the same machine) using telnet, a network protocol that allows remote access to a command-line interface. The script mostly fails because most ports are closed or not listening for connections. However, the script succeeds on ports 80 and
443, which are the default ports for HTTP and HTTPS protocols, respectively. These protocols are used for web services and web browsers. Therefore, the technician should look for an unauthorized HTTP service on this server, as it may indicate a security breach or a misconfiguration. Looking for a virus infection on this server is also possible, but not the most likely source of the issue. Looking for an unauthorized Telnet service on this server is not relevant, as the script is using telnet as a client, not a server. Looking for an unauthorized port scanning service on this server is not relevant, as the script is scanning ports on the localhost, not on other machines. References:
* https://phoenixnap.com/kb/telnet-windows
* https://www.techopedia.com/definition/23337/http-port-80
* https://www.techopedia.com/definition/23336/https-port-443

NO.3 An administrator is troubleshooting an application performance issue on a virtual server with two vCPUs. The application performance logs indicate CPU contention. The administrator adds more vCPU cores to the VM, yet the issue persists. Which of the following is the most likely reason for this issue?

A. The server has high page utilization.

B. The server has high disk latency.

C. The application is single-threaded.

D. The application cannot be virtualized.

Answer: C
Explanation:
A single-threaded application is an application that can only execute one task or process at a time. A single- threaded application can only utilize one CPU core, regardless of how many cores are available or assigned to the virtual machine. Therefore, adding more vCPU cores to the VM will not improve the performance of the application, as it will still be limited by the speed and capacity of one core12. To troubleshoot this issue, the administrator should check if the application is single-threaded or multi- threaded. This can be done by using tools such as Task Manager, Performance Monitor, or Process Explorer on Windows, or top, htop, or ps on Linux34. If the application is single-threaded, the administrator should consider the following options:
* Reduce the number of vCPU cores on the VM to match the number of threads that the application can use. This can help avoid CPU contention and co-stop issues that may arise from having too many vCPUsrelative to the number of physical cores on the host5.
* Upgrade the physical CPU on the host to a faster or newer model that can provide higher clock speed and performance for the single core that the application uses.
* Optimize the application code or configuration to make it more efficient or multi-threaded, if possible.
This can help the application take advantage of multiple cores and improve its performance.

NO.4 A company uses a hot-site, disaster-recovery model. Which of the following types of data replication is required?

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

NO.5 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.

NO.6 An administrator is patching a file and print server After rebooting the server, it begins acting strangely. The administrator tries to open Print Management, but it will not open Upon inspection, the spooler service has been disabled. The administrator then notices the server has services that should be enabled but are disabled.
Which of the following actions should the administrator take next to resolve the issue as quickly as possible?
A. Roll back the installed updates
B. Set the services to restart automatically and reboot the server
C. Enable the spooler service.
D. Reboot the server into safe mode

Answer: B
Explanation:
When a server is acting strangely after a patch and services that should be enabled are disabled, it's often a good first step to set the services to restart automatically and then reboot the server. This can resolve many issues as it ensures that all services start correctly upon reboot

NO.7 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.

NO.8 Which of the following will correctly map a script to a home directory for a user based on username?

A. \\server\users$\username

B. \\server\%username%

C. \\server\FirstInitialLastName

D. \\server\$username$

Answer: B
Explanation:
The administrator should use \server%username% to correctly map a script to a home directory for a user based on username. %username% is an environment variable that represents the current user's name on a Windows system. By using this variable in the path of the script, the administrator can dynamically map the script to the user's home directory on the server. For example, if the user's name is John, the script will be mapped to \server\John.
Reference:
https://social.technet.microsoft.com/Forums/windows/en-US/07cfcb73-796d-48aa-96a9-08280a1ef25a/ mapping-home-directory-with-username-variable?forum=w7itprogeneral

NO.9 Which of the following backup types only backs up files that have changed since the last full backup?

A. Differential

B. Open file

C. Incremental

D. Snapshot

Answer: A

Explanation:

Understanding different backup types is crucial for effective data protection strategies. Here's a breakdown of the relevant backup methods:

* Full Backup: Captures all data, regardless of previous backups.
* Differential Backup: Backs up data that has changed since the last full backup.
* Incremental Backup: Backs up data that has changed since the last backup, whether it was full or incremental.
* Snapshot: Captures the state of a system at a specific point in time.

A Differential Backup starts with a full backup. Subsequent differential backups save copies of all files that have been modified since that full backup. This means each differential backup includes all changes made since the last full backup, leading to larger backup sizes over time but faster restoration, as only the last full backup and the latest differential backup are needed.

In contrast, an Incremental Backup also begins with a full backup, but each subsequent backup only includes data that has changed since the most recent backup (whether full or incremental). This approach results in smaller backup sizes and quicker backup processes. However, restoration can be slower and more complex, as it requires the last full backup and all subsequent incremental backups to fully restore data.

Therefore, the correct answer is A. Differential, as it specifically refers to backing up files that have changed since the last full backup.

References:

CompTIA Server+ Certification Exam Objectives (SK0-005): Backup Methods CompTIA Server+ (SK0-005) Study Guide: Chapter on Security and Disaster Recovery

NO.10 A technician notices that every time a server is powered on, it turns off after several minutes. After reviewing logs, the technician notices the server registers execution of the shutdown. Which of the following should the technician do to fix this issue?

A. Reset the memory modules

B. Check for fan failure

C. Change the VM password

D. Set credentials for the remote console

Answer: B

Explanation:

The server is shutting down after being powered on, which could indicateoverheating. This is a common issue whenfans fail, causing the CPU or other components to overheat and forcing the system to shut down to protect itself from damage.

* Check for fan failure (Answer B):Ensuring the server's cooling system is functioning properly is crucial. Overheating due to a fan failure can cause the system to shut down automatically.
* Resetting memory modules (Option A):While memory issues can cause system instability, they generally do not lead to immediate shutdowns as described.
* Changing the VM password (Option C):This is unrelated to the shutdown issue.
* Setting credentials for the remote console (Option D):This is irrelevant to the described problem of server shutdown.

CompTIA Server+ Reference:This topic is related toSK0-005 Objective 3.3: Diagnose hardware and software issues.

NO.11 A certain application initially uses 1TB of drive space, but this is expected to double each year for the next two years. Which of the following is the minimum number of 1TB drives that are needed in a RAID 5 configuration?

A. 3

B. 4

C. 5

D. 6

Answer: C

Explanation:

RAID 5 is a storage configuration that uses striping with parity, providing both improved performance and fault tolerance. It requires a minimum of three disks, where data and parity information are distributed across all drives. The storage capacity of a RAID 5 array is calculated as (N - 1) * S, where N is the number of drives, and S is the size of each drive.

Storage Requirements:

* Initial Storage: 1TB
* After 1 Year: Doubles to 2TB
* After 2 Years: Doubles again to 4TB

To accommodate 4TB of data in a RAID 5 setup, we use the formula:

(N - 1) * 1TB # 4TB

Solving for N:

N - 1 # 4

N # 5

Therefore, a minimum of 5 drives, each 1TB in size, is required to meet the projected storage needs. This configuration will provide a total usable capacity of 4TB, with 1TB allocated for parity to ensure fault tolerance.

References:

CompTIA Server+ Certification Exam Objectives (SK0-005): RAID Levels and Types CompTIA Server+ (SK0-005) Study Guide: Chapter on Storage Solutions

NO.12 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/

NO.13 A server administrator wants to check the open ports on a server. Which of the following commands should the administrator use to complete the task?
A. nslookup
B. nbtstat
C. telnet
D. netstat -a
Answer: D
Explanation:
netstat is a command-line tool that displays network connections, routing tables, interface statistics, and more.
The -a option shows all listening and non-listening sockets on the server. This can help check the open ports on a server and identify any unwanted or malicious connections.References:https://docs.microsoft.com/en-us /windows-server/administration/windows-commands/netstat

NO.14 A VLAN needs to be configured within a virtual environment for a new VM. Which of the following will ensure the VM receives a correct IP address?
A. A virtual router
B. A host NIC
C. A VPN
D. A virtual switch
E. A vNIC
Answer: D
Explanation:
The correct answer is D. A virtual switch.
A virtual switch is a software-based network device that connects the virtual machines (VMs) in a virtual environment and allows them to communicate with each other and with the physical network. A virtual switch can also create and manage virtual LANs (VLANs), which are logical segments of a network that separate the traffic of different VMs or groups of VMs. A VLAN needs a DHCP server to assign IP addresses to the VMs that belong to it. A virtual switch can act as a DHCP relay agent and forward the DHCP requests from the VMs to the DHCP server on the physical network.This way, the VMs can receive correct IP addresses for their VLANs123 A virtual router is a software-based network device that routes packets between different networks or subnets.
A virtual router can also create andmanage VLANs, but it is not necessary for a VM to receive a correct IP address.A virtual router can be used to provide additional security, redundancy, or load balancing for the VMs12 A host NIC is a physical network interface card that connects the host machine to the physical network. A host NIC can also support VLAN tagging, which allows the host machine to communicate with different VLANs on the network. However, a host NIC alone cannot ensure that a VM receives a correct IP address for its VLAN.The host NIC needs to be connected to a virtual switch that can relay the DHCP requests from the VMs to the DHCP server12 A VPN is a virtual

private network that creates a secure tunnel between two or more devices over the internet.
A VPN can be used to encrypt and protect the data traffic of the VMs, but it is not related to the configuration of VLANs or IP addresses.A VPN does not affect how a VM receives a correct IP address for its VLAN14 A vNIC is a virtual network interface card that connects a VM to a virtual switch or a virtual router. A vNIC can also support VLAN tagging, which allows the VM to communicate with different VLANs on the network.
However, a vNIC alone cannot ensure that a VM receives a correct IP address for its VLAN.The vNIC needs to be connected to a virtual switch or a virtual router that can relay the DHCP requests from the VMs to the DHCP server12

NO.15 Which of the following BEST measures now much downtime an organization can tolerate Curing an unplanned outage?

A. SLA

B. BIA

C. RTO

D. MTTR

Answer: C

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

NO.16 A data center employee shows a driver's license to enter the facility Once the employee enters, the door immediately doses and locks, triggering a scale that then weighs the employee before granting access to another locked door. This is an example of.

A. mantrap.

B. a bollard

C. geofencing

D. RFID.

Answer: A

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

References:
* https://www.techopedia.com/definition/16293/mantrap
* https://www.techopedia.com/definition/1437/bollard
* https://www.techopedia.com/definition/23961/geofencing
* https://www.techopedia.com/definition/506/radio-frequency-identification-rfid

NO.17 A large number of connections to port 80 is discovered while reviewing the log files on a server. The server is not functioning as a web server. Which of the following represent the BEST immediate actions to prevent unauthorized server access? (Choose two.)

A. Audit all group privileges and permissions

B. Run a checksum tool against all the files on the server

C. Stop all unneeded services and block the ports on the firewall

D. Initialize a port scan on the server to identify open ports

E. Enable port forwarding on port 80

F. Install a NIDS on the server to prevent network intrusions

Answer: C F

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

NO.18 A recent power Outage caused email services to go down. A sever administrator also received alerts from the datacenter's UPS.
After some investigation, the server administrator learned that each POU was rated at a maximum Of 12A.
INSTRUCTIONS
Ensure power redundancy is implemented throughout each rack and UPS alarms are resolved. Ensure the maximum potential PDU consumption does not exceed 80% or 9.6A).
a. PDU selections must be changed using the pencil icon.
b. VM Hosts 1 and 2 and Mail Relay can be moved between racks.
c. Certain devices contain additional details

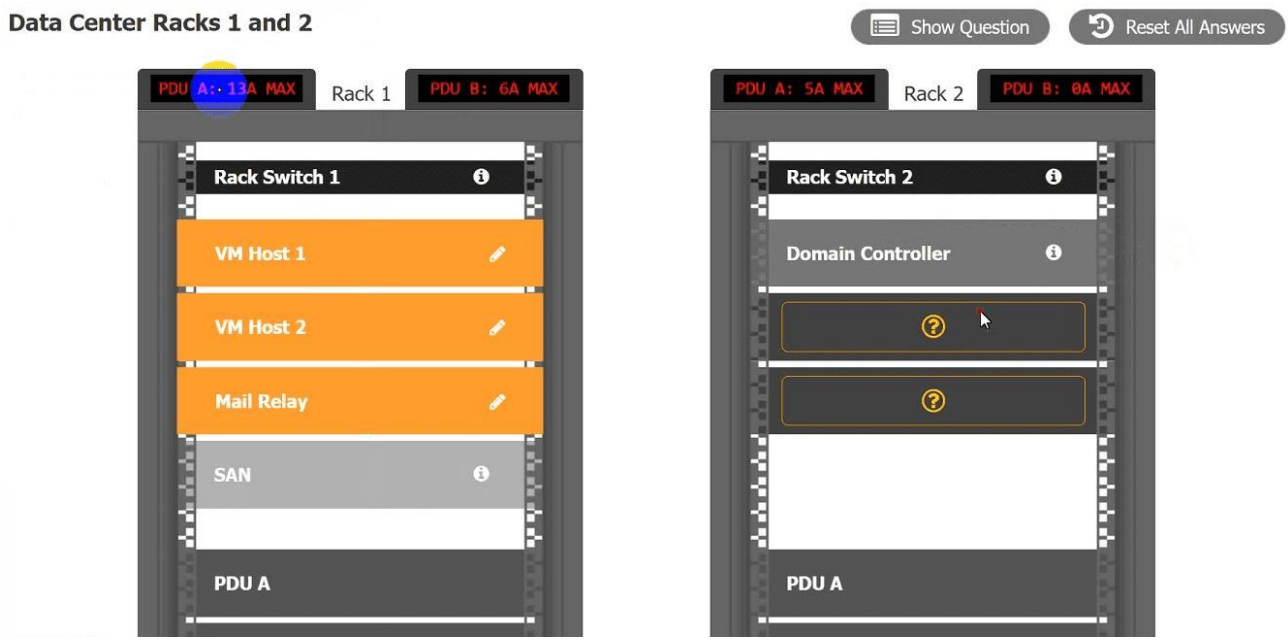

Answer:

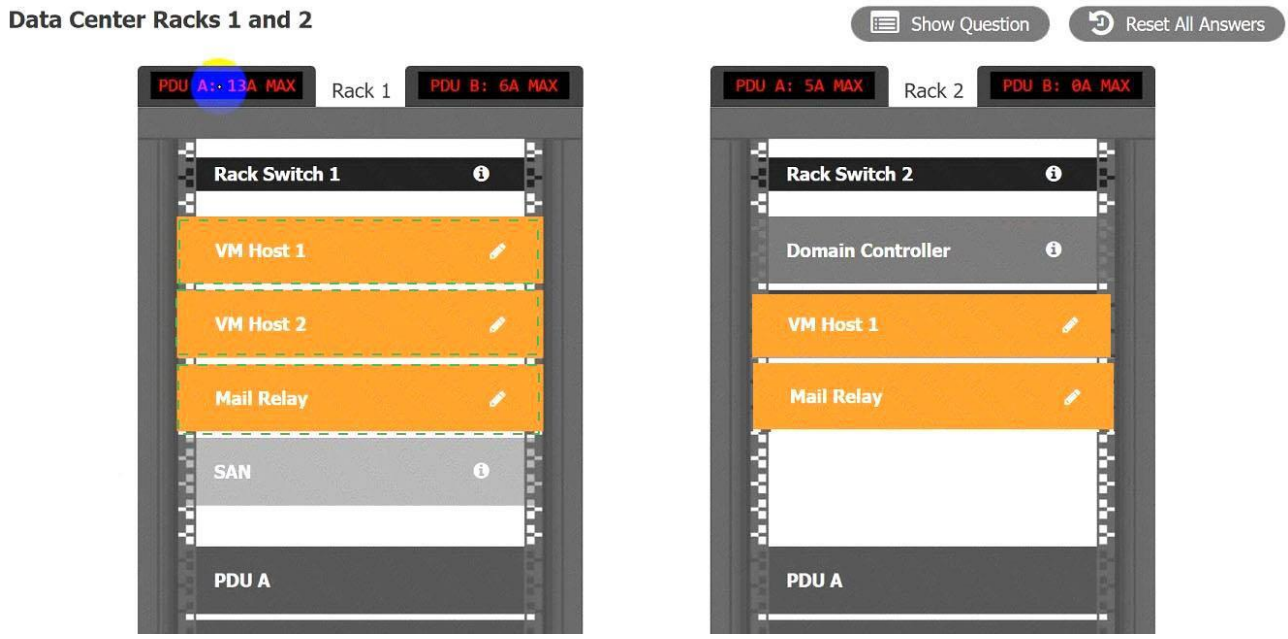

NO.19 A server administrator has received tickets from users who report the system runs very slowly and various unrelated messages pop up when they try to access an internet-facing web application using default ports. The administrator performs a scan to check for open ports and reviews the following report:
Starting Nmap 7.70 https://nmap.org) at 2019-09-19 14:30 UTC
Nmap scan report for www.abc.com (172.45.6.85)
Host is up (0.0021s latency)
Other addresses for www.abc.com (not scanned) : 4503 : F7b0 : 4293: 703: : 3209 RDNS record for 172.45.6.85: 1ga45s12-in-f1.2d100.net Port State Service
21/tcp filtered ftp
22/tcp filtered ssh
23/tcp filtered telnet
69/tcp open @username.com

80/tcp open http
110/tcp filtered pop
143/tcp filtered imap
443/tcp open https
1010/tcp open www.popup.com
3389/tcp filtered ms-abc-server
Which of the following actions should the server administrator perform on the server?

A. Close ports 69 and 1010 and rerun the scan.

B. Close ports 80 and 443 and rerun the scan.

C. Close port 3389 and rerun the scan.

D. Close all ports and rerun the scan.

Answer: A

Explanation:

Port 69 is used for TFTP (Trivial File Transfer Protocol), which is an insecure and unencrypted protocol for file transfer. Port 1010 is used for a malicious website that generates pop-up ads. Both of these ports are likely to be exploited by hackers or malware to compromise the server or the web application. The server administrator should close these ports and rerun the scan to verify that they are no longer open12.

References = 1: Why Are Some Network Ports Risky, And How Do You Secure Them? - How-To Geek (https://www.howtogeek.com/devops/why-are-some-ports-risky-and-how-do-you-secure-them/) 2: Switchport Port Security Explained With Examples - ComputerNetworkingNotes(https://www.computernetworkingnotes.com/ccna-study-guide/switchport-port-security-explained-with-examples.html)

NO.20 A systems administrator has several different types of hard drives. The administrator is setting up a MAS that will allow end users to see all the drives within the NAS. Which of the following storage types should the administrator use?

A. RAID array

B. Serial Attached SCSI

C. Solid-state drive

D. Just a bunch of disks

Answer: D

Explanation:

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

NO.21 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.

NO.22 Which of me following is the BEST action to perform before applying patches to one of the hosts in a high availability cluster?

A. Disable the heartbeat network.

B. Fallback cluster services.

C. Set the cluster to active-active.

D. Failover all VMs.

Answer: D

Explanation:

This is the best action to perform before applying patches to one of the hosts in a high availability cluster. A high availability cluster is a group of hosts thatact like a single system and provide continuous uptime. A high availability cluster is often used for load balancing, backup, and failover purposes. Failover is a process of transferring workloads from one host to another in case of a failure or maintenance. By failing over all VMs (Virtual Machines) from the host that needs to be patched to another host in the cluster, the technician can ensure that there is no downtime or data loss during the patching process. Disabling the heartbeat network is not a good action to perform, as this would disrupt the communication and synchronization between the hosts in the cluster. Fallback cluster services is not a valid term, but it may refer to restoring cluster services after a failover, which is not relevant before applying patches. Setting the cluster to active-active is not a good action to perform, as this would increase the load on both hosts and reduce redundancy. References: https://www.howtogeek.com/190014/virtualization-basics-understanding-techniques-and-fundamentals/https://www.howtogeek.com/428483/what-is-end-to-end-encryption-and-why-does-it-matter/

NO.23 Which of the following is a system that scans outgoing email for account numbers, sensitive phrases, and other forms of PII?

A. SIEM

B. DLP

C. HIDS

D. IPS

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

NO.24 An administrator is configuring a server that will host a high-performance financial application. Which of the following disk types will serve this purpose?
A. SAS SSD
B. SATA SSD
C. SAS drive with 10000rpm
D. SATA drive with 15000rpm
Answer: A
Explanation:
The best disk type for a high-performance financial application is a SAS SSD. A SAS SSD (Serial Attached SCSI Solid State Drive) is a type of storage device that uses flash memory chips to store data and has a SAS interface to connect to a server or a storage array. A SAS SSD offers high speed, low latency, high reliability, and high durability compared to other types of disks, such as SATA SSDs, SAS HDDs, or SATA HDDs. A SAS SSD can handle high I/O workloads and deliver consistent performance for applications that require fast data access and processing.
Reference:
https://www.hp.com/us-en/shop/tech-takes/sas-vs-sata

NO.25 A technician is attempting to reboot a remote physical Linux server. However, attempts to command a shutdown -----now result in the loss of the SSH connection. The server still responds to pings. Which of the following should the technician use to command a remote shutdown?
A. virtual serial console
B. A KVM
C. An IDRAC
D. A crash cart
Answer: C
Explanation:
An IDRAC (Integrated Dell Remote Access Controller) is a tool that can be used to command a remote shutdown of a physical Linux server. An IDRAC is a hardware device that provides out-of-band management for Dell servers. It allows the technician to access the server's console, power cycle, reboot, or shut down the server remotely using a web interface or a command-line interface. An IDRAC does not depend on the operating system or network connectivity of the server. A virtual serial console is a tool that can be used to access a remote virtual machine's console using a serial port connection. A KVM (Keyboard, Video, Mouse) switch is a device that allows the technician to switch between different computer sources using the same keyboard, monitor, and mouse. A crash cart is a mobile unit that contains a keyboard, monitor, mouse, and other tools that can be connected to a physical server for troubleshooting purposes. References: https://www.

dell.com/support/kbdoc/en-us/000131486/understanding-the-idrac
https://www.howtogeek.com/799968/what- is-a-kvm
-switch/https://www.techopedia.com/definition/1032/business-impact-analysis-bia

NO.26 Which of the following would MOST likely be part of the user authentication process when implementing SAML across multiple applications?

A. SSO

B. LDAP

C. TACACS

D. MFA

Answer: A

Explanation:

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

Reference:

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

NO.27 A server is only able to connect to a gigabit switch at 100Mb. Other devices are able to access the network port at full gigabit speeds, and when the server is brought to another location, it is able to connect at full gigabit speed. Which of the following should an administrator check first?

A. The switch management

B. The VLAN configuration

C. The network cable

D. The network drivers

Answer: C

Explanation:

The first thing that the administrator should check is the network cable. The network cable is a physical medium that connects a server to a switch or other network device. The network cable can affect the speed and quality of the network connection, depending on its type, length, and condition. If the network cable is damaged, faulty, or incompatible, it can cause the server to connect at a lower speed than expected. Therefore, the administrator should check the network cable for any signs of wear, tear, or mismatch, and replace it if necessary.

References: CompTIA Server+ SK0-005 Certification Study Guide, Chapter 2, Lesson 2.1, Objective 2.1

NO.28 A server administrator wants to run a performance monitor for optimal system utilization. Which of the following metrics can the administrator use for monitoring? (Choose two.)

A. Memory

B. Page file

C. Services

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

NO.29 A server technician installs a new NIC on a server and configures the NIC for IP connectivity. The technician then tests the connection using thepingcommand. Given the following partial output of thepingandipconfigcommands:

```
ipconfig /all

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

pinging 192.168.1.1 with 32 bytes of data:

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

Which of the following caused the issue?
A. Duplicate IP address
B. Incorrect default gateway
C. DHCP misconfiguration
D. Incorrect routing table
Answer: A
* The ping command output shows that the NIC has an IP address of 192.168.1.100 and a default gateway of 192.168.1.1. However, when the technician tries to ping the default gateway, the reply comes from another IP address: 192.168.1.101. This means that there is another device on the network that has the same IP address as the default gateway, and it is responding to the ping request instead of the intended destination.
* A duplicate IP address can cause network connectivity problems, such as packet loss, routing errors, or unreachable hosts. To resolve this issue, the technician should either change the IP address of the default gateway or the device that is conflicting with it, or use DHCP to assign IP addresses automatically and avoid conflicts.

* The other options are not correct because they do not explain the ping output. An incorrect default gateway would cause no reply or a destination unreachable message, not a reply from a different IP address. A DHCP misconfiguration would cause an invalid or no IP address on the NIC, not a duplicate IP address on the network. An incorrect routing table would cause routing errors or unreachable destinations, not a reply from a different IP address.
References:
https://askleo.com/what_is_ping_and_what_does_its_output_tell_me/
https://learn.microsoft.com/en-us/windows-server/administration/windows-commands/ping

NO.30 Which of the following documents would explain the consequences of server downtime?

A. Service-level agreement

B. Business continuity plan

C. Disaster recovery plan

D. Business impact analysis

Answer: D

Explanation:
Abusiness impact analysis (BIA)is a document that outlines the potential effects of downtime on business operations. It identifies critical business functions and estimates the impact of disruption on the organization's ability to function.
* Business impact analysis (Answer D):This document is specifically designed to assess the consequences of downtime and other disruptions.
* Service-level agreement (Option A):Defines the expected level of service between a provider and a client but doesn't directly explain downtime consequences.
* Business continuity plan (Option B):Outlines how the business will continue operating during a disruption but doesn't focus solely on the consequences of downtime.
* Disaster recovery plan (Option C):Focuses on restoring systems after a disaster but doesn't outline the specific impact of downtime.
CompTIA Server+ Reference:This topic relates toSK0-005 Objective 4.1: Explain disaster recovery concepts.

NO.31 The management team has mandated the use of data-at-rest encryption on all corporate servers. Using this encryption paradigm will ensure:

A. website traffic is protected while traversing the internet.

B. files stored on the server are protected against physical theft.

C. attachments that are emailed from this server cannot be intercepted.

D. databases in use are protected from remote hackers.

Answer: B

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

NO.32 A server administrator is currently working on an incident. Which of the following steps should the administrator perform before resolving the issue?

A. Inform the impacted users.
B. Make the changes to the system.
C. Determine the probable causes.
D. Identify changes to the server.
Answer: C
Explanation:
The step that the server administrator should perform before resolving the issue is to determine the probable causes. This step is part of the troubleshooting process that follows a logical and systematic approach to identify and solve problems with servers and applications. The troubleshooting process consists of several steps, such as:
* Identify the problem: Gather information from various sources, such as users, logs, or alerts, to understand the symptoms and scope of the problem.
* Establish a theory of probable cause: Analyze the information and formulate one or more possible causes of the problem based on evidence or experience.
* Test the theory to determine cause: Perform tests or experiments to verify or eliminate each possible cause until the root cause is found.
* Establish a plan of action to resolve the problem and implement the solution: Design and execute a plan to fix the problem using appropriate tools and techniques.
* Verify full system functionality and implement preventive measures: Confirm that the problem is resolved and that no other issues arise as a result of the solution. Implement preventive measures to avoid recurrence of the problem or improve performance.
* Document findings, actions, and outcomes: Record the details of the problem, its cause, its solution, and its outcome for future reference or knowledge sharing. References: [CompTIA Server+ Certification Exam Objectives], Domain 6.0: Troubleshooting, Objective 6.1: Given a scenario involving server hardware issues (e.g., power supply failure), troubleshoot using appropriate tools.

NO.33 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

NO.34 A technician is preparing a deployment of servers to be used by staff at a remote location.

Which of the following should the technician do to prevent access to the hardware configuration?

A. Enable an administrator account

B. Enable a UEFI password

C. Disable WOL

D. Enable encryption at rest

Answer: B

Explanation:

Enabling aUEFI (Unified Extensible Firmware Interface) passwordprevents unauthorized users from making changes to the server's hardware configuration settings, such as boot order or device settings. This is crucial for protecting the integrity of the server at a remote location where physical security might be more difficult to enforce.

* UEFI password (Answer B):It provides security at the firmware level, preventing changes to low-level configurations unless the correct password is provided.
* Administrator account (Option A):While important for OS-level access, it doesn't prevent someone with physical access from altering hardware settings via UEFI/BIOS.
* Disabling WOL (Option C):Wake-on-LAN (WOL) allows a device to be powered on remotely. Disabling it can help with security but does not prevent hardware configuration changes.
* Encryption at rest (Option D):Encryption protects data on the server but does not prevent hardware configuration access.

CompTIA Server+ Reference:This topic is covered underSK0-005 Objective 2.1: Install and configure server operating systems.

NO.35 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.