

Practice Exam Questions

CompTIA



CompTIA Server+ Certification



EXAMKILLER

Help Pass Your Exam At First Try

CompTIA

Exam SK0-005

CompTIA Server+ Certification Exam

Version: 7.0

[Total Questions: 210]

Question No : 1

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

- A. Network connectivity on the server
- B. LVM status on the server
- C. Disk space in the /var directory
- D. YUM dependencies

Answer: C

Explanation:

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

Question No : 2

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

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

Answer: B

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

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

Question No : 3 HOTSPOT

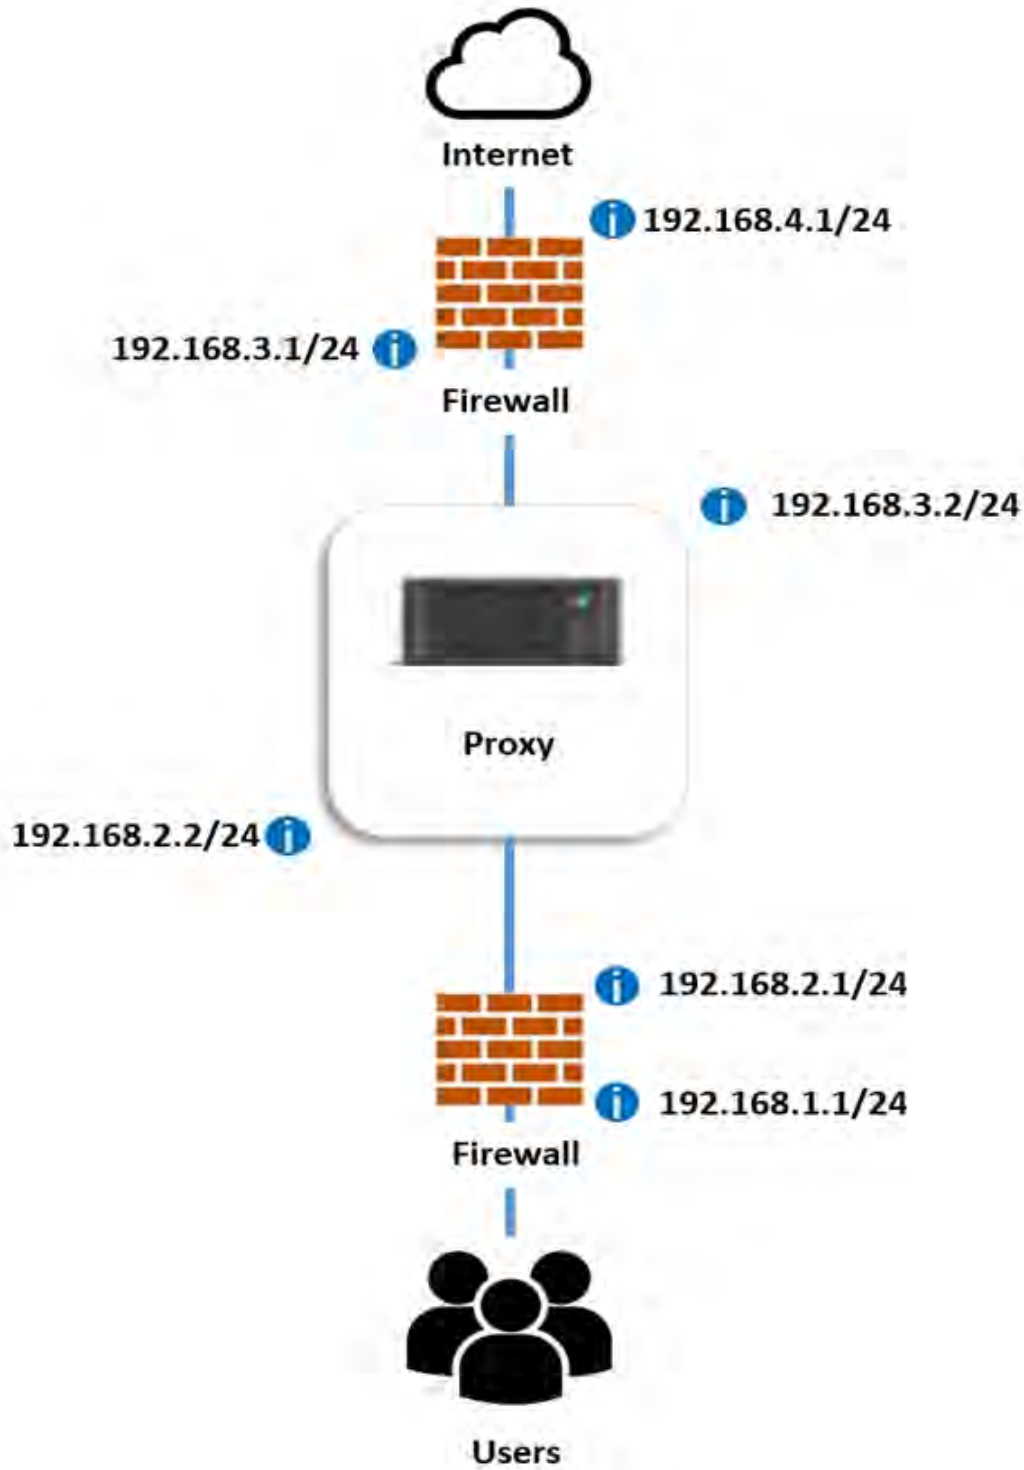
A systems administrator deployed a new web proxy server onto the network. The proxy server has two interfaces: the first is connected to an internal corporate firewall, and the second is connected to an internet-facing firewall. Many users at the company are reporting they are unable to access the Internet 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

Answer:

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

Explanation:

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

Table

Description automatically generated

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

Table

Description automatically generated

Question No : 4

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

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

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

Answer: A,E

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

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

Question No : 5

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

NEXT to install the firmware?

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

Answer: B

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

Question No : 6

Which of the following ensures a secondary network path is available if the primary connection fails?

- A. Link aggregation
- B. Most recently used
- C. Heartbeat
- D. Fault tolerance

Answer: D

Explanation: Fault tolerance is the ability of a system to continue functioning in the event of a failure of one or more of its components. Fault tolerance can ensure a secondary network path is available if the primary connection fails. Fault tolerance can be achieved by using redundant components, such as network cards, cables, switches, routers, etc., that can take over the function of the failed component without interrupting the service.

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

Question No : 7

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

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

Answer: B

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

may not work if the SSH connection is lost. RDP (Remote Desktop Protocol) is a protocol that allows remote access to a Windows server's graphical user interface, but it does not work on Linux servers and requires network connectivity. References:

<https://www.techopedia.com/definition/13623/crash-cart>

<https://www.techopedia.com/definition/13624/kvm-switch>

<https://www.techopedia.com/definition/3422/remote-desktop-protocol-rdp>

Question No : 8

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

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

Answer: B

Explanation:

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

Question No : 9

A security manager is concerned that a rogue employee could boot a server from an outside USB drive. Which of the following actions can be taken to reduce this risk? (Select TWO).

- A. Close unneeded ports.
- B. Disable unneeded physical ports.
- C. Set a BIOS password.
- D. Install a SIEM.

- E. Disable unneeded services.
- F. Install a HIDS.

Answer: B,C

Explanation: Disabling unneeded physical ports would prevent unauthorized devices from being connected to the server, such as an outside USB drive. Setting a BIOS password would restrict access to the boot settings and prevent unauthorized changes to the boot order. The other options would not address the risk of booting from an outside USB drive

Question No : 10

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

unauthorized server access? (Choose two.)

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

Answer: C,F

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

Question No : 11

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>

Question No : 12

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

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

Answer: A

Explanation: Open-source software is software that is freely available and can be modified and distributed by anyone. It usually requires very little technical support and has no licensing fees. Therefore, it would be the lowest cost solution for an application that does not need much support. References: <https://www.comptia.org/training/resources/exam-objectives/comptia-server-sk0-005-exam-objectives> (Objective 2.3)

Question No : 13

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

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

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

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

Answer: B,D,E

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

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

References:

- ✍ <https://phoenixnap.com/kb/raid-levels-and-types>
- ✍ https://en.wikipedia.org/wiki/Standard_RAID_levels

Question No : 14

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

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

Answer: D,E

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

QUESTION NO: 63 Correct Answer: A Explanation: The RAID level that provides the highest possible capacity for a storage array is RAID 0. RAID 0 is a type of RAID level that provides performance enhancement by using striping. Striping means dividing data into blocks and distributing them across multiple disks to increase speed and capacity. RAID 0 does not provide any fault tolerance or redundancy, as it does not use any parity or mirroring techniques. RAID 0 uses all of the available disk space for data storage, without losing any space for overhead. Therefore, RAID 0 provides the highest possible capacity for a storage array, but also has the highest risk of data loss.

Question No : 15

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

```
Bash ./startup.sh:Permission denied
```

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

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

Answer: C

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

References: <https://www.howtogeek.com/437958/how-to-use-the-chmod-command-on-linux>

Question No : 16

A company is running an application on a file server. A security scan reports the application has a known vulnerability. Which of the following would be the company's BEST course of action?

- A. Upgrade the application package
- B. Tighten the rules on the firewall

- C. Install antivirus software
- D. Patch the server OS

Answer: A

Explanation: The best course of action for the company is to upgrade the application package to fix the known vulnerability. A vulnerability is a weakness or flaw in an application that can be exploited by an attacker to compromise the security or functionality of the system. Upgrading the application package means installing a newer version of the application that has patched or resolved the vulnerability. This way, the company can prevent potential attacks that may exploit the vulnerability and cause damage or loss.

Question No : 17

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

- A. Use an MFA out-of-band solution.
- B. Use a Telnet connection.
- C. Use a password complexity policy.
- D. Use a role-based access policy.

Answer: A

Explanation: This is the most secure method to access servers located in remote branch offices because MFA stands for multi-factor authentication, which requires users to provide more than one piece of evidence to prove their identity. An out-of-band solution means that one of the factors is delivered through a separate channel, such as a phone call, a text message, or an email. This adds an extra layer of security and prevents unauthorized access even if a password is compromised. References: <https://docs.microsoft.com/en-us/azure/active-directory/authentication/concept-mfa-howitworks>

Question No : 18

A server technician is configuring the IP address on a newly installed server. The documented configuration specifies using an IP address of 10.20.10.15 and a default gateway of 10.20.10.254. Which of the following subnet masks would be appropriate for this setup?

- A. 255.255.255.0
- B. 255.255.255.128
- C. 255.255.255.240
- D. 255.255.255.254

Answer: A

Explanation:

The administrator should use a subnet mask of 255.255.255.0 for this setup. A subnet mask is a binary number that defines how many bits of an IP address are used for the network portion and how many bits are used for the host portion. The network portion identifies the specific network that the IP address belongs to, while the host portion identifies the specific device within that network. The subnet mask is usually written in dotted decimal notation, where each octet represents eight bits of the binary number. A 1 in the binary number means that the corresponding bit in the IP address is part of the network portion, while a 0 means that it is part of the host portion. For example, a subnet mask of 255.255.255.0 means that the first 24 bits (three octets) of the IP address are used for the network portion and the last 8 bits (one octet) are used for the host portion. This subnet mask allows up to 254 hosts per network ($2^8 - 2$). In this case, the IP address of 10.20.10.15 and the default gateway of 10.20.10.254 belong to the same network of 10.20.10.0/24 (where /24 indicates the number of bits used for the network portion), which can be defined by using a subnet mask of 255.255.255.0.

Question No : 19

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

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

Answer: A

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

risk of data loss.

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

Question No : 20

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

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

Answer: C

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

Reference:

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

Question No : 21

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 login. To join a server to a domain, the administrator must have a valid domain account with sufficient privileges and must know the name of the domain controller that hosts Active Directory.

Question No : 22

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

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

Answer: C,E

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

from outside. Reflective glass can enhance privacy and security for datacenter operations.

Question No : 23

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

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

Answer: D

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

Reference:

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

Question No : 24

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

># Enforcing

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

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

Answer: A

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

Reference:

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

Question No : 25

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.

Question No : 26

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

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

Answer: D

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

Question No : 27

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

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

Answer: A

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

Question No : 28

A storage administrator is investigating an issue with a failed hard drive. A technician replaced the drive in the storage array; however, there is still an issue with the logical volume. Which of the following best describes the NEXT step that should be completed to restore the volume?

- A. Initialize the volume
- B. Format the volume
- C. Replace the volume
- D. Rebuild the volume

Answer: D

Explanation: The administrator should rebuild the volume to restore it after replacing the failed hard drive. A volume is a logical unit of storage that can span across multiple physical disks. A volume can be configured with different levels of RAID (Redundant Array of Independent Disks) to provide fault tolerance and performance enhancement. When a hard drive in a RAID volume fails, the data on that drive can be reconstructed from the remaining drives using parity or mirroring techniques. However, this process requires a new hard drive to replace the failed one and a rebuild operation to copy the data from the existing drives to the new one. Rebuilding a volume can take a long time depending on the size and speed of the drives and the RAID level.

Question No : 29

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/>

Question No : 30

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: A synthetic full backup is a backup method that describes copying only changed files each time it runs and using only a single file during a restore. A synthetic full backup is a backup approach that involves creating a new full backup by using the previous full backup and related incremental backups. This means that a backup solution does not have to transfer the full amount of data from the source machine and can synthesize the latest incremental backups with the last full backup to create a new full backup. This reduces the backup window and network bandwidth consumption. During a restore, only the latest synthetic full backup file is needed to recover the data. Open file backup is a backup method that allows backing up files that are in use or locked by applications. Full incremental backup is a backup method that involves performing a full backup first and then backing up only the changed files since the last backup. Full differential backup is a backup method that involves performing a full backup first and then backing up only the