

CompTIA Cloud+ (CV0-004) Version 2.0.1 - Quiz Questions with Answers

1.0 Cloud Architecture

1.0 Cloud Architecture

1.

Acme Inc. is looking for low-cost connectivity to public cloud providers. They are willing to accept some unpredictability in performance to keep costs low.

What solution is the BEST fit for these requirements?

Public internet

Using a dedicated connection

Renting bandwidth from other organizations using the same public cloud

Building their own global network

Correct answer: Public internet

The Internet is a shared global network with enormous bandwidth and capacity. Organizations can implement their own connections to a public cloud provider, but that is more expensive.

Organizations typically do not allow other organizations to use their bandwidth for security reasons.

2.

In which of the following cloud service models is the cloud service provider responsible for the data center hardware running the hypervisors?

IaaS, PaaS, and SaaS

IaaS and PaaS only

PaaS only

IaaS only

Correct answer: IaaS, PaaS, and SaaS

The cloud service provider is responsible for hardware required to run systems that deliver cloud services in the SaaS, PaaS, and IaaS models.

3.

What type of cloud model is a pool of computing resources that is delivered over the Internet by a third-party provider and used by many organizations?

Public

Hybrid

Private

Community

Correct answer: Public

The public cloud is a set of computing services offered by third-party providers over the public Internet that are available to any organization that wants to use or purchase those services.

Private clouds are dedicated to one organization.

Hybrid cloud describes the use of both public and private clouds.

Community cloud is a cloud dedicated to a specific group of organizations with a shared objective and similar requirements.

4.

What level of drive redundancy does RAID 0 provide?

None

N+1

N+N

N+2

Correct answer: None

Redundant array of independent disks (RAID) 0 does not provide redundancy. If a drive fails in a RAID 0 configuration, data will be lost.

Therefore, all the "N+" answers are incorrect.

5.

You do not want to incur the cost of storing all your data in the most expensive storage because, over time, data becomes less frequently accessed. Typically, once data is over a year old, it is accessed less than twice per year.

What storage provisioning strategy should you use to strike a balance between performance, accessibility, and cost?

Storage tiers

RAID classes

Caching

Thick provisioning

Correct answer: Storage tiers

Storage tiers allow you to combine different classes of storage with different cost structures into a single storage pool. Data that is requested often may be stored in higher, more expensive, higher-speed storage. Data that is requested less frequently can be stored in less expensive tiers.

RAID classes would help define different levels of fault tolerance, but would not directly address the case of optimizing storage access based on usage.

Caching is useful for frequently accessed data, but would not directly address the use case of infrequently accessed data.

Thick provisioning allocates all the defined storage capacity to a given drive upon creation.

6.

What technique do firewalls use to determine if a session has been established for an incoming packet before accepting it?

Stateful packet inspection

Session packet inspection

Firewall packet inspection

NAT packet inspection

Correct answer: Stateful packet inspection

Stateful packet inspection monitors the state of active connections and uses this information to determine which network packets to allow through the firewall.

Firewall packet inspection is a generic term that can generally refer to any type of packet inspection performed by a firewall.

Session packet inspection and NAT packet inspection are not standard terms a Cloud+ candidate needs to be familiar with.

7.

What address space is required for a Class A private network?

10.0.0.0 to 10.255.255.255

172.16.0.0 to 172.31.255.255

172.16.0.0 to 172.18.255.255

192.0.0.0 to 192.168.0.0

Correct answer: 10.0.0.0 to 10.255.255.255

Private network address spaces are commonly used in an internal corporate network. They are not made available as public addresses. The Class A private network address space of 10.0.0.0 to 10.255.255.255 generates over 16 million addresses.

The other answers listed are related to the class B or C ranges.

8.

Which of the following metrics is directly related to compute resources on a virtual machine?

Wait time

Paged pool

Page faults

Virtual NIC average bytes received

Correct answer: Wait time

Compute resources deal with compute technologies like CPU, GPU, and microprocessors.

Wait time is a CPU-related metric a CPU thread waits for processing.

Paged pool and page faults are related to memory.

Virtual NIC average bytes received is related to network resources.

9.

Which of the following levels of service availability offers the MOST uptime?

Five nines

Four nines

99.9%

90.999%

Correct answer: Five nines

Five nines is a common term that means 99.999% availability.

Similarly, four nines represents 99.99% availability.

99.9% and 90.999% are both lower levels of service availability.

10.

An organization's employees need to access the Internet from their internal devices through the company firewall.

What mechanism is used to hide the internal device IP address from the outside world?

Network Address Translation (NAT)

Port Address Masking (PAM)

Address Masking Translation (AMT)

Internet Address Masking (IAM)

Correct answer: Network Address Translation (NAT)

NAT enables private IP networks to connect to the Internet by translating the internal addresses to one public IP address. Only one address is exposed to the outside world. This provides security by effectively hiding the entire internal network behind that single address.

The other answers are distractor acronyms, not standard protocols, technologies, or techniques.

11.

You wish to monitor resources in order to understand which users consume which resources and for how long, in order to allocate costs back to departments so that they pay for their share of compute resource usage.

What is this policy called?

Chargeback

Show back

Cost split

Cost share

Correct answer: Chargeback

Instead of simply charging all IT costs to one central department, an organization charges costs to individual users or departments. This is known as a chargeback.

Show back, cost split, and cost share are distractor answers.

12.

You are supporting a global application in the cloud that is accessed by employees across 20 countries. The application has a service level agreement (SLA) mandating a user response time of less than one second. Results of a recent connectivity test show that all users were able to access the application; however, users in three countries experienced average response times of five to ten seconds.

Would this be considered a successful connectivity test?

No, because successful connectivity also means latency objectives must be met, especially if there are SLAs.

Yes, because all users were able to connect to the application successfully.

Yes, because the majority of users were able to connect to the application and experienced a good response time.

Yes, because the organization did not experience any downtime.

Correct answer: No, because successful connectivity also means latency objectives must be met, especially if there are SLAs.

Connectivity to a cloud service is defined as both the physical connection to cloud resources and acceptable latency (delay). Slow cloud performance due to connectivity issues can have business impacts on organizations. In this example, there was a published SLA for a sub-second response time. Because a population of users did not experience sub-second response times, this would be considered an unsuccessful connectivity test.

13.

Terraform and Ansible are examples of what type of platform?

Orchestration

Community

Monitoring

Observability

Correct answer: Orchestration

Terraform and Ansible are both orchestration platforms. Orchestration platforms are tools for automating the resource and infrastructure management.

14.

Which of the following is an example of increased read/write storage speed?

IOPS increasing from 60 to 80

Threads decreasing from 10 to 1

Disk utilization increasing from 40% to 90%

RAM increasing from 8GB to 16GB

Correct answer: IOPS increasing from 60 to 80

Input/output operations per second (IOPS) is a measurement of reads or writes from storage per second. Higher IOPS numbers indicate increased read/write storage speed.

Threads are related to CPUs, not storage.

Disk utilization increasing from 40% to 90% would not necessarily increase read/write storage speed.

RAM is memory, not storage. Increasing RAM would not directly increase read/write storage speed, although it can improve performance and may indirectly influence IOPS numbers.

15.

Acme Inc. follows standard practices for encryption at rest and encrypts data at rest on a server that uses solid-state drives (SSDs) and has 128 GB of RAM.

When would the data on the server be encrypted?

Before it is written to the SSD

Before it is written to RAM

After it is written to the SSD

After it is written to RAM

Correct answer: Before it is written to the SSD

Data at rest is encrypted before it is written to storage. SSDs and hard drive disks (HDDs) are examples of storage.

RAM is memory, not storage.

16.

Which type of service level agreement (SLA) is the MOST complex to test for to determine if the cloud provider has met their contractual service obligation?

Multilevel SLA

Customer-based SLA

Service-based SLA

Cloud-based SLA

Correct answer: Multilevel SLA

A service level agreement (SLA) is a contract that outlines the type of service to be provided, the service's performance level, how the performance levels will be monitored and reported on, steps for reporting service issues, issue resolution process, timeframe, and service provider penalties if the SLA is not met. Multilevel SLAs are the most complicated because they are created for different types of consumers who use the same services.

A service-based SLA describes a single service that is provided for all customers.

A customer-based SLA is differentiated between business customers and home users.

There is no type of SLA known as a cloud-based SLA.

17.

Cruz is a database administrator at Acme Inc. who regularly searches a MySQL database to provide insights to management and internal stakeholders.

Cruz wants to improve the performance of commonly searched data in the database. This data is not frequently written.

What should Cruz do to achieve his goal?

Index the data

Change the datatypes

Migrate to a Redis database

Use functions

Correct answer: Index the data

Indexing data in a relational database can significantly improve search performance. Indexing can negatively impact write speeds because additional data is associated with indexed data.

The datatypes are not given in the question, and, generally, changing datatypes alone does not meaningfully improve search speed.

Migrating a relational MySQL database to a nonrelational Redis database is a complex task that may not improve search performance at all.

Functions are database operations such as sum or count. Nothing in the question suggests they would directly impact search performance.

18.

In which cloud service model does the cloud provider have the MOST responsibility for security, compliance, and liability requirements?

SaaS

PaaS

IaaS

XaaS

Correct answer: SaaS

Software as a Service (SaaS) providers are responsible for maintaining security levels for the user; the user has no way of configuring the security settings. Security, compliance, and liability requirements are all built into the SaaS application.

SaaS shifts more burden to the cloud service provider than PaaS or IaaS. The tradeoff for users is the level of control.

XaaS is a generic term that means "anything as a service" and is not a specific model.

19.

A video streaming service uses an algorithm to suggest videos you may like based on your previous viewing patterns. The suggestions update over time based on the videos you view and "like."

This is MOST likely an implementation of what type of technology?

ML

MTU

REST

Serverless

Correct answer: ML

Suggestion algorithms are a common use case for machine learning (ML).

MTU is a network setting.

REST is a standard pattern for creating and using APIs.

Serverless computing is a form of computing where all the underlying resources are abstracted away from the consumer and functions are executed ad-hoc.

20.

Charlie is a database administrator at Acme Inc. Charlie takes a full backup of a SQL database server by saving the entire database as a collection of files.

What term BEST describes this type of database backup?

Database dump

Volume backup

LUN cloning

File backup

Correct answer: Database dump

A database dump is a full database backup that saves an entire database in one or more files.

Volume backup is a generic term that can refer to backing up any storage volume.

File backups are backups of filesystems that are often used for workstations and servers.

LUN cloning creates a copy of a specific unit of block storage.

21.

A content delivery network (CDN) is an implementation of what type of secondary data store?

Replica

Synthetic

Differential

Cloud NAS

Correct answer: Replica

Replicas improve reliability and fault tolerance because backup copies are replicated to multiple storage devices at different locations. A CDN is an example of replicas in practice. CDNs are commonly used to distribute copies of data such as video and web assets for better performance.

Network attached storage (NAS) is a type of file storage commonly used for sharing files with multiple users. A NAS is not necessarily a secondary data store.

Synthetic is a term that could be part of multiple concepts, such as synthetic testing or synthetic full backups (which are full backups made from an old full backup and incremental backups).

Differential is a backup strategy that involves backing up changes since the last full backup.

22.

Which of the following is a prerequisite for detecting an anomaly with cloud infrastructure?

Baseline

Syslog

IaC

CPU monitoring

Correct answer: Baseline

Anomalies represent a deviation from a baseline value. Therefore, without a baseline, you cannot have an anomaly.

Syslog is a logging protocol. While it could be used to detect some anomalies, it is not a requirement.

CPU monitoring is required for detecting CPU anomalies, but there are other types of cloud infrastructure anomalies (e.g., RAM, I/O, network utilization, etc.) that do not require CPU monitoring.

Infrastructure as code (IaC) is an approach to maintaining infrastructure configuration like source code.

23.

You are a Software as a Service (SaaS) cloud user. Which of the following application components is NOT provided by your SaaS cloud provider?

Physical client device

User interface

Data storage

Program

Correct answer: Physical client device

An application consists of a user interface, data, and the program.

Cloud providers do not provide a physical device to gain access to the platform. There could be a wide variety of choices for the end user, including mobile devices and laptops.

24.

Acme Inc. has a legacy accounting application that currently runs on an on-premises server. They are replacing the application by decoupling its components and deploying them across multiple containers in a cluster to create a new application.

What type of architecture is the new application using?

Microservices

Cloud

Monolith

MVC

Correct answer: Microservices

Microservices architecture involves breaking an application down into smaller, typically containerized, components that are loosely coupled and serve a specific purpose in a larger application.

Cloud is too generic of a term to be applicable to this question.

Monolith architectures include all the services and components needed to run an application deployed on the same system. The legacy application is an example of a monolith.

Model-view-controller (MVC) is an approach to web application design.

25.

Which type of storage is attached to the computer and does not require network conductivity to be accessed by the computer?

Direct-attached storage (DAS)

Network-attached storage (NAS)

Object storage

Hardware file share (HFS)

Correct answer: Direct-attached storage (DAS)

DAS is physically connected to the host computer through internal hardware.

HFS is not a standard type of storage.

Object storage is a type of storage for unstructured data.

Network-attached storage (NAS) is accessed over a network.

26.

Which of the following statements about VXLAN and GENEVE is FALSE?

VXLAN is VMware's proprietary implementation of the open-source GENEVE standard

VXLAN and GENEVE both allow transport of Ethernet traffic over an IP network

VXLAN and GENEVE can both help address size limitations with traditional Ethernet

VXLAN and GENEVE use different default ports

Correct answer: VXLAN is VMware's proprietary implementation of the open-source GENEVE standard

Virtual extensible LAN (VXLAN) is a network virtualization technique that encapsulates Ethernet frames in IP/UDP packets. It was created to address size limitations associated with Ethernet.

Generic network virtualization encapsulation (GENEVE) is an alternative solution for the same basic problem VXLAN aims to solve.

A key difference between GENEVE and VXLAN is that GENEVE only defines a packet format but does not prescribe how packets are sent.

VXLAN uses port 4789 by default, while GENEVE uses port 6081 by default.

VXLAN is not a VMware proprietary implementation.

27.

Acme Enterprises runs an accounting firm in the European Union (EU). Which of the following could impede an organization's ability to replicate data across multiple datacenters on different continents?

Regulatory requirements

N + 1 rules

EIA/TIA 568

VT-x

Correct answer: Regulatory requirements

Regulatory requirements can restrict where data is stored. For example, in the EU, GDPR places data residency restrictions on the personal data of EU citizens.

N + 1 is a type of redundancy that can support a single system failure without downtime.

EIA/TIA 568 is a structured cabling standard.

VT-x is a virtualization technology for x86 CPUs.

28.

Your organization is running out of storage because users have stored many redundant copies of their data.

What technology would you use to free up space?

Deduplication

Caching

Garbage collection

Compression

Correct answer: Deduplication

Deduplication removes redundant data from a storage system by finding the location of the copies and placing a pointer to those copies instead of storing the physical data multiple times.

Compression could decrease the size of the pieces of redundant data, which would improve storage utilization. However, deduplication is the better answer because it removes the redundant copies completely.

Caching can improve read performance on stored data, but would not directly address duplicate data and storage space challenges.

Garbage collection is a memory management technique used by many programming languages.

29.

You have configured a Linux virtual server to act as a DHCP server on a virtual private network. You configured the server to use the default DHCP port.

Client devices cannot receive the DHCP addresses from the server. During troubleshooting, you confirm that the client requests are not reaching the DHCP server.

Unblocking which of the following ports is MOST likely to solve the problem?

68

53

123

989

Correct answer: 68

68 is the default DHCP port.

989 and 990 are the default FTPS ports.

123 is the default NTP port.

53 is the default DNS port.

30.

You are a cloud network administrator at Acme Inc. You enable support for jumbo frames on several network appliances.

Which of the following will likely happen as a result?

Less processor utilization on the networking appliances

Increased TTLs on the networking appliances

Decreased TTLs on the networking appliances

Increased hop counts on the networking appliances

Correct answer: Less processor utilization on the networking appliances

Jumbo frames allow network devices to support packet payloads up to 9,000 bytes. Supporting larger packet payloads leads to less network-related processor utilization than with smaller payload sizes.

Enabling jumbo frames would not directly change network TTLs or hops.

31.

You are a developer at Acme Inc. You want to modify what features are enabled or disabled when a container is running based on the system it is deployed on.

What can you use to achieve the desired outcome?

Environment variables

Hardcoding

RDP

NVRAM

Correct answer: Environment variables

Environment variables can store values that change how a container works depending on what the variable is set to. Different runtime environments can set different environment variables that can cause the same container image to have different functionality when it is instantiated.

Hardcoding is effectively the opposite of using environment variables and involves writing code that includes specific values.

Remote desktop protocol (RDP) is a remote access protocol.

NVRAM is a type of memory that can persist its state after being powered off.

32.

Which of the following BEST describes a use case where an organization may use network segmentation?

To optimize security and performance

To reduce network costs

To maximize available network addresses

To combine micro networks

Correct answer: To optimize security and performance

Network segmentation divides a network into small pieces called segments. This allows the organization to apply specific policies and rules to the segments. The more segments there are, the harder it is for attackers to spread their malicious code. Segments can also have different levels of resources applied to them, depending on their performance needs.

Combining smaller networks is the opposite of segmenting them.

Network segmentation does not directly reduce network costs or increase available network addresses.

33.

A solid-state drive is connected to a physical server using a SCSI interface. This is an example of which type of storage?

DAS

NAS

SATA

iSATA

Correct answer: DAS

Direct-attached storage (DAS) is storage that is connected via a physical connection to a server such as a SCSI or SATA interface. Cloud providers sometimes use DAS for ephemeral storage.

Network-attached storage (NAS) is storage that is accessible over a network. NAS typically uses LAN connectivity.

SATA (serial AT attachment) connections are a way storage can be directly attached.

iSATA is a distractor answer.

34.

What is a major reason scripting languages are used for orchestration?

Simple to learn

Large footprint

Limited support

Require many updates

Correct answer: Simple to learn

Scripting languages use commands that are English-like and intuitive, thereby making them easy to learn and also easily maintainable.

None of the other answers are an advantage that would make scripting languages popular for orchestration tasks.

35.

A smart sensor that connects to a network and streams humidity data to a public cloud is an example of what type of technology?

IoT

CSP

MTU

MQTT

Correct answer: IoT

Internet of Things (IoT) devices are a category of network-connected (typically Internet-connected) devices that serve some primary purpose other than traditional computing. Common examples of IoT devices include "smart" sensors, cameras, lights, and other embedded systems.

A cloud service provider (CSP) provides cloud services to cloud customers.

Maximum transmission unit (MTU) is a network setting.

MQTT is a publish/subscribe style messaging protocol often used by IoT devices.

36.

You are a developer creating several containers for an application that will run on cloud-based virtual machines.

Which of the following statements about your deployment is TRUE?

The container and host kernels will be shared

You are building a monolithic application

You cannot use Windows to run your containers

The container and host IP addresses must match

Correct answer: The container and host kernels will be shared

A container shares a kernel with its host operating system (OS).

Containers are not typically associated with monolithic applications. Containers are more commonly associated with microservices architecture.

Windows systems can run containers in general. There is nothing in the question that says the containers cannot be Windows containers.

Container and host IP addresses do not need to match.

37.

Acme Cloud LLC. is a cloud service provider. They provide an accounting software to cloud consumers. The software is accessible to authorized users over the Internet.

Acme Inc. maintains all the hardware, hypervisors, operating systems, databases, and patching required to deliver the service. Cloud consumers are responsible for user account management.

What cloud service model is Acme Cloud LLC. using?

SaaS

IaaS

Private

Public

Correct answer: SaaS

Software as a service (SaaS) is a cloud service model where the consumer has the least responsibility. The cloud service provider is responsible for hardware, patching, hypervisors, databases, operating systems, and more. The consumer is typically responsible for application data and user account management with SaaS.

In the infrastructure as a service (IaaS) model, the consumer would be responsible for databases and operating systems.

Public and private cloud are cloud deployment models, not cloud service models.

38.

You have a server with 64GB of RAM which runs a type 1 hypervisor.

You have 10 virtual machines with requirements for 8 GB of RAM each. The virtual machines typically only use 1-2 GB of RAM in practice.

What is the most resource-efficient way you can deploy all 10 virtual machines?

Use overcommitting and deploy all 10 VMs on the hypervisor

Purchase a second server and run 5 VMs on each server

Convert the hypervisor to a type 2 hypervisor

Purchase a server for each VM

Correct answer: Use overcommitting and deploy all 10 VMs on the hypervisor

Overcommitting allows virtual machines to have more memory than is physically available allocated to them. Overcommitting memory is typically based on the assumption that VMs will not use all the RAM allocated to them at the same time.

Purchasing additional servers is less resource-efficient than using overcommitting.

A type 2 hypervisor is typically less resource-efficient than a type 1 hypervisor.

39.

What type of redundancy allows a cluster to withstand the complete loss of a system?

N + 1

Blue-Green

Five nines

MPLS

Correct answer: N + 1

N + 1 redundancy is used to describe a cluster that can withstand the loss of a single system without failure.

Blue-Green is a software deployment model.

Five nines is a term used to describe 99.999% uptime.

Multiprotocol label switching (MPLS) is a networking technique.

40.

You are a technology consultant for Acme Inc. Acme Inc. wants to determine which use cases for AI/ML make the most sense for their organization.

Which of the following use cases is LEAST likely to benefit from AI/ML?

A text editor for Windows users

A system to predict delivery arrival times

A facial recognition application

A fraud detection system

Correct answer: A text editor for Windows users

Artificial intelligence (AI) and machine learning (ML) are useful for use cases that involve making predictions, pattern recognition (e.g., facial recognition), and anomaly recognition (e.g., fraud detection).

Because a text editor is a straightforward and deterministic application, it is less likely to benefit from AI/ML than a prediction system, facial recognition system, or fraud detection system.

41.

How does a virtual private network (VPN) provide a secure connection to the cloud?

By sending traffic through an encrypted tunnel

By sending encapsulated traffic across the LAN

By only sending data to trusted endpoints

By using onion routing

Correct answer: By sending traffic through an encrypted tunnel

A VPN connection establishes a secure connection between the user and a public communications channel such as the Internet. All data traffic transmitted by the user through the Internet to the cloud is sent through an encrypted virtual tunnel. The user's IP address is masked, making its identity anonymous to others on the network.

VPNs do not necessarily only connect to trusted endpoints.

Sending encapsulated traffic across a LAN would not enable the outbound connections over an untrusted network (like the Internet) VPN connections typically facilitate.

Onion routing is a routing technique that sends traffic through a network of decentralized routers.

42.

Acme Inc. wants to replicate data between two datacenters in different regions. To connect the datacenters, they want to use a high-speed fiber optic connection that is not owned or operated by a telco.

What connectivity solution is the BEST choice for Acme Inc.'s requirements?

Dark fiber

MPLS

VPN

EIA/TIA 568

Correct answer: Dark fiber

Dark fiber is fiber optic cable connections that are not owned or operated by the telcos that maintain many of the connection mediums modern enterprises lease.

MPLS lines are often operated by telcos. MPLS networks are not necessarily connected using fiber.

VPNs are encrypted network tunnels that connect network endpoints.

EIA/TIA 568 is a structured cabling standard.

43.

What do routers use to determine the most efficient path to send packets to their destination?

Routing tables

DMZ

Port tables

Load tables

Correct answer: Routing tables

A routing table is a data table stored in a router that contains information about the topology of the network immediately around it and the next appropriate "hop" for network traffic.

A demilitarized zone (DMZ) is a specific portion of a network that sits between a trusted internal network and the public Internet.

Port tables and load tables are distractor answers.

44.

You are purchasing physical servers that will act as hypervisors in your organization's private cloud. The servers have x86 Intel processors.

What technology will enable enhanced hardware virtualization support for these servers?

VT-x

AMD-V

EIA

vNIC

Correct answer: VT-x

VT-x is an Intel technology that enables enhanced hardware virtualization support for select Intel processors.

AMD-V is a similar technology for AMD processors.

EIA standards are commonly associated with cabling.

vNICs are virtualized network interface cards.

45.

Integer, date, and decimal are examples of what type of database element?

Data types

Schemas

Functions

Views

Correct answer: Data types

Integer, date, and decimal are all examples of data types.

Database schemas define how a database is structured.

Database functions are computations such as count or sum.

Database views are alternative ways to view the data in a database table(s).

46.

Amal, a QA engineer at Acme Inc., is tasked with testing a cloud application against a specification that details the business logic the application must follow.

What type of testing will Amal perform?

Functional

Usability

Performance

Regression

Correct answer: Functional

Functional testing focuses on testing how an application should (or should not) perform. It is typically based on a specification that defines the expected behavior.

Usability testing focuses on how easy to use ("usable") a system is.

Performance testing focuses on a system's ability to perform under load.

Regression testing focuses on testing for issues created by changes or updates.

In this case, Amal is tasked with conducting functional tests to validate that the cloud application is working as expected.

47.

An organization's employees need to access the Internet from their internal devices through the company firewall.

What technique can the organization use to map multiple internal IP addresses to a single public IP address?

Port Address Translation (PAT)

Media access control (MAC)

Router Address Translation (RAT)

Internet Address Translation (IAT)

Correct answer: Port Address Translation (PAT)

PAT is used to translate private IP addresses into a public IP address via port numbers. Translating addresses from multiple devices to a single public IP address is a common use case for PAT.

Media access control (MAC) is a layer 2 networking concept that does not deal with mapping private and public IP addresses.

Internet Address Translation (IAT) and Router Address Translation (RAT) are distractor answers.

48.

Single-root I/O virtualization (SR-IOV) allows multiple virtual machines (VMs) to share a single instance of what physical component?

NIC

RAM

CPU

HDD

Correct answer: NIC

SR-IOV allows multiple VMs to use a single physical NIC.

SR-IOV does not directly impact RAM, CPU, or HDD allocation.

49.

Acme Inc. has a "five nines" uptime SLA with their customers. Which of the following is the HIGHEST amount of downtime allowed per year that would not breach the SLA?

4 minutes

40 minutes

400 minutes

None of the answers are within the SLA

Correct answer: 4 minutes

A "five nines" (99.999%) service level agreement (SLA) for uptime would allow for 5.39 minutes of downtime per year. 4 minutes of downtime per year would not breach a "five nines" SLA.

50.

You have configured a Linux virtual server to act as a DNS server on a virtual private network. You configured the server to use the default DNS port. Client devices cannot query the DNS server.

Unblocking which of the following ports is MOST likely to solve the problem?

53

443

68

123

Correct answer: 53

DNS uses TCP and UDP port 53 by default.

- *443 is the default HTTPS port.*
 - *123 is the default NTP port.*
 - *68 is the default DHCP port.*
-

51.

You are a developer at Acme Inc. You would like to run multiple versions of the same application on your local Linux workstation without the applications interfering with one another.

What is the MOST lightweight (fewest resources) way to achieve this goal?

Containers

Virtual machines

Community cloud

Blue-green deployments

Correct answer: Containers

Running multiple containers on the same Linux system is a way to run multiple versions of an application without using as many resources as virtual machines would require.

A community cloud is a cloud deployment model.

Blue-green deployments is a software deployment model where one instance of an application is serving production users while another is used for updates and testing. Once updates and testing complete, the updated instance begins to serve production users, and the other instance is updated and tested.

52.

Which CPU technology performs virtualization functions directly in the computer's hardware instead of software?

VT-x

Hyperthreading

Ballooning

Physical virtualization

Correct answer: VT-x

VT-x runs a hypervisor's virtualization instructions directly in the hardware instead of running them in the software. This makes the hypervisor perform faster.

Hyperthreading allows a single core to be represented as separate CPUs. It depends on the underlying operating system or hypervisor supporting multiprocessing.

"Physical virtualization" is a distractor term.

Ballooning is a technique that enables unused memory to be reclaimed.

53.

After you have deployed your virtual machines in the cloud, users complain about slow performance.

What metric will tell you if the cause is likely virtual machines waiting to access a CPU for processing?

CPU wait time

Page faults

Average bytes sent and received

CPU pool size

Correct answer: CPU wait time

CPU wait time shows the amount of time the CPU is processing workloads. If the CPU time is running consistently high from the beginning, this is an indicator that it was not configured correctly because there is little to no room for growth.

A page fault is an exception related to memory management.

Average bytes sent and received is related to network traffic.

The CPU pool size deals with the allocation of virtual CPUs. It would not directly tell you if the virtual machines were waiting for CPU resources.

54.

Jumbo frames allow support for Ethernet frame payloads of up to what size?

9,000 bytes

1,500 bytes

1 MB

1 GB

Correct answer: 9,000 bytes

Jumbo frames support Ethernet frame payloads of up to 9,000 bytes. Enabling jumbo frames can reduce network device processor utilization.

55.

You are running a high-availability cluster that rebalances the virtual machines for performance. A specific application will benefit from using processor cache consistently. Your hypervisor's automatic CPU thread allocation has not adequately addressed this requirement,

What approach can you use to make sure that the application's VM consistently uses the same CPU and can gain the performance benefits of using the CPU cache?

CPU affinity

Dedicated rebalancing

Reserved virtual machines

High-performance pooling

Correct answer: CPU affinity

CPU affinity ties processes from a specific virtual machine to a specific type of CPU. Typically, it is best to allow the hypervisor to automatically allocate CPU threads. However, a primary use case for CPU affinity is to optimize CPU cache-related performance.

Dedicated rebalancing, reserved virtual machines, and high-performance pooling would not directly address the CPU cache-related challenge in the question.

56.

What type of cloud cost management model tracks cloud usage metrics in order to bill individual departments for their actual cloud usage?

Chargeback

Showback

Usage-based

Depreciation

Correct answer: Chargeback

A chargeback model tracks cloud usage by departments, organizations, or other entities and bills them for their actual usage.

Showback is a distractor answer.

Usage-based is a generic term that can describe any billing model based on usage.

Depreciation is an asset's loss of value over time.

57.

Your global customers require a 24 x 7 x 365 availability service-level agreement (SLA). That means none of your data centers can afford to experience any failure.

What is the BEST replication strategy for your business requirements?

Multi-regional replication

Statewide replication

Caching replication

Semi-regional replication

Correct answer: Multi-regional replication

Multi-regional replication duplicates data across multiple geographies. This strategy offers the highest availability because if there is a failure in one site, there are several other sites that are configured to immediately take over the workload with no interruption.

Statewide replication and semi-regional replication are both distractor answers that would not be as fault tolerant as multi-regional replication.

Caching replication is a distractor answer that would only address availability for cached resources (for example, data stored on a cloud delivery network).

58.

Which of the following statements related to containerization is TRUE?

Environment variables are not encrypted

Windows systems cannot run Docker containers

Linux systems cannot run OCI-compliant containers

Container kernels are isolated from host kernels

Correct answer: Environment variables are not encrypted

Environment variables are not encrypted and should not be used for sensitive data. Secrets should be used instead.

Windows can run Docker containers.

Linux can run containers, including containers that are compliant with OCI runtime specifications.

The kernel of the host operating system is shared with the containers that run on it.

59.

You have set up multiple servers in the cloud to equally distribute the processing of your user workload. You have recently noticed that the utilization of three of the servers is different from the remaining servers.

What type of testing would you do FIRST to understand why this is happening?

Load balancing testing

Compute testing

Storage testing

Connectivity testing

Correct answer: Load balancing testing

The purpose of a load balancer is to intercept incoming traffic and equally distribute the workload to the back-end servers. When the load balancer is functioning properly, all of the servers receive equal workloads. Therefore, their utilization rates are similar to each other. In this scenario, the network is functioning properly, so we direct our attention to the load balancer and run tests. If it is determined that the load balancer is functioning properly, then you do additional testing on the individual servers' hardware components.

Compute testing would focus on processor (CPU) resources.

Storage testing would focus on disk-related issues.

Connectivity testing focuses on network connectivity. While a load balancer is a network component, connectivity is working in general in the scenario described.

60.

A tester at Acme Inc. is trying to guess the username and password to access a web server in a cloud network under test. They are using common default credentials and easily guessable passwords. Their attempts fail, and they are unable to authenticate using their brute force attempt.

What type of testing is the tester MOST likely performing?

Penetration

Vulnerability

Regression

Functional

Correct answer: Penetration

Penetration testing focuses on attempting to exploit vulnerabilities in a system, network, or infrastructure. Penetration tests typically involve a tester outside of the system (e.g., from another network or without credentials) attempting to gain access. Trying easy-to-guess passwords is a common penetration testing technique.

Vulnerability testing focuses on detecting vulnerabilities and security threats that might be exploited. Vulnerability testing typically involves the use of a vulnerability scanner that compares an environment to a database of known vulnerabilities. Since the tester is actively trying to exploit a possible weakness and there was no mention of a scanning tool, penetration testing is a better answer than vulnerability testing.

Regression testing focuses on testing for issues created by changes or updates. There is no indicator that this testing is being performed after a change or update.

Functional testing focuses on testing how an application should (or should not) perform. It is typically based on a specification that defines the expected behavior. The question does not suggest that the tester is testing the system for expected functionality.

61.

Which of the following is NOT a benefit of nonrelational databases when compared to relational databases?

Data deduplication

Flexible data organization

Broad data type support

Simple design

Correct answer: Data deduplication

Compared to relational databases, nonrelational databases are more likely to be impacted by data duplication that increases storage requirements.

Benefits of nonrelational databases compared to relational databases include flexible data organization, broad data type support, and simpler design.

62.

Alex is an IT administrator for a small accounting firm. The firm needs to share files between eight users within the LAN, each with their own workstation. Alex implements a solution that uses a single storage server and the SMB v3 protocol to make the files available to all the users. There is no dedicated storage network in the configuration.

Which technology did Alex use to implement file sharing in this case?

NAS

DAS

SAN

Object storage

Correct answer: NAS

NAS (network attached storage) involves a server providing access to files over a TCP/IP network instead of connecting directly to a computer. NAS implementations use standard protocols such as SMB (server message block) or NFS (network file system).

DAS (direct attached storage) involves connecting devices directly to storage drives.

SAN (storage attach networks) decouples storage and services and implements networks dedicated for storage traffic. SAN enables high performance and uses network protocols such as Fibre Channel, NVMeOF, or iSCSI.

Cloud object storage is a popular approach to cloud storage. It typically uses HTTPS or a cloud provider API instead of a protocol like NFS or SMB to enable storage access.

63.

Which of the following statements about PPTP is TRUE?

It is deprecated

It provides TLS encryption natively

It provides SSL encryption natively

It cannot encapsulate layer 2 PPP frames

Correct answer: It is deprecated

The point-to-point tunneling protocol (PPTP) is a deprecated protocol that encapsulates layer 2 point-to-point (PPP) frames. It does not provide encryption natively.

64.

Acme Inc. wants to have most of the benefits of a dedicated private cloud, but use public cloud infrastructure.

Which cloud model is the BEST fit for this scenario?

Cloud within a cloud

Community cloud

Hybrid cloud

SaaS

Correct answer: Cloud within a cloud

The cloud within a cloud model enables organizations to have the flexibility a private cloud offers while using a public cloud provider's infrastructure. This provides the organization with a dedicated set of resources without the need for an on-premises server often required for a private cloud.

A community cloud is a cloud used by multiple organizations with a shared interest and similar requirements.

A hybrid cloud is a mix of public and private cloud usage. The reason hybrid cloud is not correct in this case is that a hybrid cloud does not necessarily mean that dedicated resources will be provisioned for the customer.

SaaS offers application access in a cloud service model.

65.

Running containers typically use what type of storage?

Nonpersistent

Persistent

DAS

OCI-compliant MPLS

Correct answer: Nonpersistent

Containers typically run in nonpersistent storage (e.g., memory).

Direct attached storage refers to storage that is connected physically using an interface like SATA or SCSI.

OCI-compliant MPLS is a distractor answer. MPLS is a networking technique, not a type of storage.

66.

Which of the following statements about availability zones (AZs) and regions is TRUE?

Regions typically consist of multiple AZs

AZs and regions typically have a 1:1 relationship

AZs typically consist of multiple regions

AWS supports regions but not availability zones

Correct answer: Regions typically consist of multiple AZs

A region typically consists of two or more availability zones (AZs).

AWS recognizes both the AZ and region concepts.

67.

Where is persistent storage for containerized applications in a microservices architecture typically stored?

Outside of the container

In DAS

Inside a dedicated storage container

Inside each container in a microservices architecture

Correct answer: Outside of the container

Persistent storage is typically provided outside of (not inside) a container. This is useful because containers typically run in nonpersistent storage and are often intended to be "stateless."

Direct attached storage (DAS) is storage that is directly connected using a physical connection such as SCSI or SATA.

68.

You have configured a Linux virtual server to act as an FTPS server on a virtual private network. You configured the server to use the default FTPS port.

Client devices cannot connect to the FTPS server. During troubleshooting, you confirm that the client requests are not reaching the FTPS server.

Unblocking which of the following ports is MOST likely to solve the problem?

989 and 990

53 and 68

20 and 21

25 and 123

Correct answer: 989 and 990

989 and 990 are the default FTPS ports.

- *20 and 21 are default FTP (not FTPS) ports.*
 - *123 is the default NTP port.*
 - *53 is the default DNS port.*
 - *25 is the default SMTP port.*
 - *68 is the default DHCP port.*
-

69.

Acme Inc. is experiencing lower than desired network throughput from a specific branch office. They ask you to recommend options that could increase network throughput.

Which of the following is a viable method for increasing network throughput?

Bonding network adapters

Increase network hops

Increase packet TTLs

Turn of page faults

Correct answer: Bonding network adapters

Bonding network adapters can significantly increase network throughput.

Increasing network hops (the number of networks a packet crosses before its destination) or packet time-to-live (TTL) settings would not increase network throughput directly.

Page faults is a metric related to memory, not network throughput.

70.

You are a cloud architect responsible for allocating resources. What approach should you take for testing sizing changes?

Add resources independently and test each one for performance impact before scaling further

Add all resources at once, test as a group for performance impact, and un-allocate as necessary

Allocate memory and compute resources together, test for performance impact, and scale up or down as needed

There is no need to test for sizing changes since you only pay for what you use in the cloud

Correct answer: Add resources independently and test each one for performance impact before scaling further

Testing best practice is to add resources one by one incrementally, and test after each one to measure the impact.

If all resources or more than one are added at the same time, it is very difficult to isolate which resource had which impact.

Even though cloud is a pay-as-you-go model, adding resources does incur costs.

71.

Your organization needs to extend the on-premises network into the public cloud.

What type of protocol uses a combination of encapsulation and encryption to create a secure connection over the Internet?

Tunneling

Dynamic Network Address Translation

Encrypted mapping

Secure hypertext

Correct answer: Tunneling

Tunneling uses encapsulation and encryption to create a secure connection between devices to emulate local network conductivity. Encapsulation packages data within another package to the next layer of security.

Dynamic Network Address Translation is used to map internal IP addresses to external IP addresses.

Encrypted mapping is a distractor answer and is not a standard networking protocol or technique.

Secure hypertext is a distractor answer based on the popular hypertext transfer protocol secure (HTTPS) protocol.

72.

Which of the following statements about IaaS, SaaS, and PaaS is TRUE?

SaaS places the least amount of responsibility on the consumer

IaaS places the least amount of responsibility on the consumer

PaaS places the least amount of responsibility on the consumer

The consumer responsibility in SaaS, PaaS, and IaaS is the same

Correct answer: SaaS places the least amount of responsibility on the consumer

The shared responsibility model defines the split in responsibilities between the cloud service provider and the consumer of cloud services.

Software as a service (SaaS) has fewer responsibilities for the consumer than the platform as a service (PaaS) and infrastructure as a service (IaaS) models.

73.

What encapsulation method does Stateless Transport Tunneling (STT) use?

TCP/IP

UDP

GRE

SFTP

Correct answer: TCP/IP

Stateless transport tunneling (STT) uses TCP/IP encapsulation and is primarily supported by VMware/Broadcom.

VXLAN uses UDP for encapsulation.

NVGRE uses GRE for encapsulation.

SFTP is a layer 7 protocol that uses SSH to support encrypted file transfers.

74.

VXLAN is designed to address size limitations of which older standard?

Ethernet

802.11g

GENEVE

DMZ

Correct answer: Ethernet

Virtual extensible LAN (VXLAN) is a network virtualization technique that encapsulates Ethernet frames in IP/UDP packets. It was created to address size limitations associated with Ethernet.

802.11g is a Wi-Fi protocol.

Generic network virtualization encapsulation (GENEVE) is an alternative solution for the same basic problem VXLAN aims to solve.

A demilitarized zone (DMZ) is a specific network segment that sits between an untrusted network (typically the Internet) and a trusted network.

75.

You have configured a storage area network (SAN) as part of a failover cluster. The clusters are connected using network interface cards (NICs) and the iSCSI protocol.

What technology on the NIC can reduce processing load on the CPUs and potentially improve SAN reliability and availability?

TCP offloading

MTU

LUN duplication

LUN deduplication

Correct answer: TCP offloading

TCP offloading is a feature available on select NICs that can reduce processing load in a SAN connected using iSCSI.

Maximum transmission units (MTU) defines the largest size of a packet a network device will allow.

Logical unit numbers (LUNs) are identifiers for storage resources. LUN duplication and deduplication are not standard terms.

76.

Acme Cloud Storage LLC. caps storage resources at 50GB per user in their basic tier.

Bola is an Acme Cloud Storage LLC. customer that purchased basic tier storage services.

Bola has already saved 50GB of data in their account. Bola uploads a 3GB video file to their Acme Cloud Storage LLC. storage account.

The platform saves the video in the storage account and sends Bola a warning message that they have exceeded their allowed storage amount.

The 50GB cap in this scenario is an example of what?

Soft limit

Hard limit

Hard quota

Soft quota

Correct answer: Soft limit

Cloud provider limits are a minimum or maximum for resource utilization. A soft limit allows the user to complete their action and use the resources but typically issues a warning or notification if it is exceeded.

A hard limit does not let the user complete their action and use the resources if it is exceeded.

A quota defines how a host allocates resources to guests.

77.

You are a network administrator at Acme Inc. You are troubleshooting a VXLAN issue on a virtual network. You realize that there is a default DENY rule in a firewall that is blocking the traffic.

Which port should you create an ALLOW rule for to allow VXLAN traffic on the default port?

4789

8080

53

3297

Correct answer: 4789

Virtual extensible LAN (VXLAN) is a network virtualization technique that encapsulates Ethernet frames in IP/UDP packets. It was created to address size limitations associated with Ethernet. It uses port 4789 by default.

None of the other numbers are default VXLAN port numbers.

78.

Acme Inc. is defining their data storage requirements. They have three categories of data:

- Group A: Frequently access data and must support high-performance read/write speeds
- Group B: Semi-frequently accessed data that does not require high-performance read/write speeds
- Group C: Archived read-only data that is rarely accessed and can have slow read speeds

Acme Inc. wants to place each data group in a different storage tier.

What storage tier would you recommend for group A data?

Correct answer: 1

There are three common storage tiers Cloud+ candidates should be familiar with:

- **Tier 1-** Highest performance, most reliable and durable. Suitable for frequently accessed and critical data.
- **Tier 2-** Lower performance than tier one. Typically used for data that does not have high read/write performance requirements.
- **Tier 3-** Rarely accessed data such as backups and archives. Typically has the slowest performance.

If all three groups of data must be in different storage tiers, tier 1 is most appropriate for group A data.

79.

Which network component in IaaS distributes incoming traffic evenly within the cloud instance in order to deliver optimum performance and high availability?

Load balancer

DNS server

Router

Firewall

Correct answer: Load balancer

A load balancer distributes cloud traffic across the Infrastructure-as-a-Service (IaaS) resources in the cloud, such as servers and storage pools, and is used to increase the number of concurrent users and reliability of the application in the cloud.

DNS servers are used for name resolution.

Routers are used to send traffic between networks.

Firewalls are primarily used for allowing or denying network traffic.

80.

Which of the following would NOT typically be stored as a container secret?

Feature settings

Password

Database connection string

Authentication keys

Correct answer: Feature settings

Sensitive data such as passwords, database connection strings, and authentication keys are good examples of data that should be stored in secrets when working with containers.

A feature setting is typically better defined in an environment variable that can be modified if needed. For example, to enable or disable a feature.

81.

Which of the following is NOT included in the service design package (SDP) documentation?

On-call support employee list

Support processes

Customer service level agreements

Technical solutions

Correct answer: On-call support employee list

The goal of the service design package (SDP) is to document the technology and processes that will be used to meet customer service level agreements (SLAs) across the organization.

Information regarding specific employees is not relevant nor required for the SDP. Employee information is appropriate for supporting scheduling and staffing, not design.

82.

Alex is a junior developer at Acme Inc. Alex wrote a Python service that they packaged as a Docker container. Alex stored an authentication key for the container to use as an environment variable on a host system so it would not be hardcoded in the container.

A senior developer flagged an issue with Alex's deployment. Alex is on their way to discuss the issue.

What problem did the senior developer MOST likely find with Alex's code or deployment?

The authentication key should not be an environment variable

Python services cannot be containerized

Alex should have hardcoded the authentication key

Docker does not support containerization of Python services

Correct answer: The authentication key should not be an environment variable

Sensitive data such as authentication keys should be stored in secrets, not environment variables or hardcoded.

Python services can be containerized, including in Docker containers.

83.

The "time-memory trade-off" concept from computer science is BEST summarized by which of the following statements?

Processing time and storage utilization have an inverse relationship

RAM utilization and storage utilization must be balanced

Read speed is proportional to write speed

Throughput and processing scale exponentially

Correct answer: Processing time and storage utilization have an inverse relationship

The "time-memory trade-off" explains that processing time and storage utilization have an inverse relationship. That is, as storage utilization goes down, processing time goes up (assuming all other variables remain the same).

This concept is important to remember when working with compression and deduplication techniques that can help optimize storage utilization.

All the other answers are distractor answers.

84.

Acme Inc. wants to determine the cost of all the cloud resources used as part of a development environment. To do this, they add a label of "env:dev" to all the relevant resources and filter for those resources in a cloud billing console.

The "env:dev" label is an example of what?

Tag

KPI

Chargeback

SLA

Correct answer: Tag

A tag is a label that includes a name and, optionally, a value. "env:dev" is an example of a tag. Tags are often used for organizing and filtering resources.

Key performance indicators (KPIs) are metrics that can be measured which track performance related to a specific goal or set of goals.

Chargebacks enable organizations to bill specific entities (e.g., people, teams, business units, departments, etc.) for their share of resource utilization.

A service level agreement (SLA) is a contractual agreement to meet a specific standard for a service. For example, SLAs often define availability and uptime requirements.

85.

Geographical dispersion is a cloud architecture and design approach that is intended to mitigate and lower what aspect of an organization's cloud deployment?

Risk

Budget

Operational stability

Scalability

Correct answer: Risk

Single points of failure or concentration of assets close to each other can create risk for an organization. Geographical dispersion reduces risk by lowering the impact of an outage or failure on an entire base of users and employees.

Geographical distribution enables techniques such as spreading workloads across multiple availability zones (AZs) and fault tolerance in the event a single region goes offline. Therefore, it is typically more expensive (so it does not lower budget) and improves reliability (so it does not lower operational stability).

Geographical distribution can improve (not lower) overall scalability.

86.

Acme Inc. wants to bill each department in the organization for its share of private cloud utilization.

What can they implement to achieve this goal?

Chargebacks

KPIs

OKRs

Smart billing

Correct answer: Chargebacks

Chargebacks enable organizations to bill specific entities (e.g., people, teams, business units, departments, etc.) for their share of resource utilization.

Key performance indicators (KPIs) are metrics that can be measured which track performance related to a specific goal or set of goals.

Objectives and key results (OKRs) are specific goals, typically in an organization or team.

Smart billing is a generic term that is often used in the electrical power industry.

87.

Acme Media Services LLC. provides video-on-demand services to customers around the globe. Currently, Acme Inc. only hosts the videos on virtual servers that reside in the United States East region on a public cloud service provider.

Customers in Australia are experiencing high latency and poor performance when streaming video.

What solution could Acme Media Services LLC. implement that would be MOST likely to improve performance for the customers in Australia?

CDN

VPN

DMZ

VT-x

Correct answer: CDN

A content delivery network (CDN) stores copies of data such as video, pictures, and static website content at edge locations that are typically closer to end users than origin servers. This can reduce latency and improve performance.

VPN creates an encrypted tunnel over another network. It would not typically reduce latency and improve performance.

A demilitarized zone (DMZ) is a specific network segment that sits between an untrusted network (typically the Internet) and a trusted network.

VT-x is a CPU virtualization technology.

88.

Which of the following public cloud service models runs applications that are primarily used by end users?

SaaS

PaaS

IaaS

XaaS

Correct answer: SaaS

End users are business users; they use the cloud to access business applications for their daily operations. The Software as a Service (SaaS) cloud model hosts business applications.

Platform as a service (PaaS) and infrastructure as a service (IaaS) are more commonly used by technical staff, such as developers.

XaaS is a generic term that means "anything as a service."

89.

You need to download files from a Linux virtual server in the cloud that is running SFTP on the default port.

What port should you ask the cloud security team to ALLOW for this transfer to be successful?

22

989

990

443

Correct answer: 22

SFTP uses TCP port 22 for file transfers.

989 and 990 are the default ports for FTPS.

443 is the default HTTPS port.

90.

An Acme Inc. developer needs to ensure that containers that run key services for an application can use a specific database connection string.

What should the developer do to securely provide these values to the container?

Secrets

Environment variables

Hardcoding

YAML

Correct answer: Secrets

Secrets are used to provide containers with secure access to sensitive data such as database connection strings and credentials.

Hardcoding sensitive data or placing it in environment variables are both less secure than using secrets.

YAML is a semi-structured data format.

91.

Which of the following statements about containers and virtual machines is TRUE?

Virtual machines provide stronger isolation than containers

Containers provide stronger isolation than virtual machines

Virtual machines are typically smaller than containers

Virtual machines follow the OCI runtime specification; containers do not

Correct answer: Virtual machines provide stronger isolation than containers

Virtual machines provide stronger isolation than containers. Containers running on the same system all share the same host operating system kernel. Virtual machines running on the same host do not share the same operating system kernel.

Containers are typically smaller than virtual machines.

The Open Container Initiative (OCI) runtime specification is relevant to containers, not to virtual machines.

92.

What is the metric used for testing storage read and write speed?

Input/output operations per second (IOPS)

Bytes of data in storage (BpS)

Bits of data in storage (bpS)

Hard drive throughput

Correct answer: Input/output operations per second (IOPS)

Input operations read data from storage and output operations write data to storage. IOPS is a performance indicator used to determine how well the storage system is performing and its utilization. IOPS testing and comparison against the baseline enables an organization to know if storage is performing as expected or if changes are needed.

Bits and bytes of data give an indication of the amount of data in storage, but not performance directly.

Hard drive throughput is a generic term and not a specific metric.

93.

What is the purpose of a QA environment?

QA is where software testers evaluate code and the potential impact on other parts of the system

QA is where programmers test their small units of code

QA is where the business end users test their code

QA is where maintenance and support take place

Correct answer: QA is where software testers evaluate code and the potential impact on other parts of the system

Deployments are implemented in three landscapes: development, quality assurance (QA), and production. Developers do their individual programming and testing in the development landscape. When they are individually done and combine their code into a function, it is ready to test as an integrated unit, and this takes place in QA.

After extensive testing and user acceptance (which takes place in a separate user environment), the code is placed into the production landscape.

Production is where maintenance and support take place.

94.

Your organization runs completely on a public cloud accessible via the public Internet.

What type of testing should you do to make sure that your cloud is reachable by end users at a given point in time?

Connectivity testing

High Availability (HA) testing

Performance testing

Security testing

Correct answer: Connectivity testing

The most basic cloud capability is to be able to connect user devices to the network. Connectivity testing checks to be sure that the network links are working and latency (delay) is minimized. If connectivity to the cloud is interrupted, the entire organization will experience downtime and is likely to incur business losses.

Nothing in the question states HA is implemented or must be tested.

Performance testing goes beyond basic connectivity.

While availability is an aspect of security, connectivity testing is more relevant to the question.

95.

Acme Inc. wants to use a cloud service model that provides them with control over the operating system that runs on their virtual machines.

Which cloud service model should they use?

IaaS

SaaS

Public

Private

Correct answer: IaaS

In the infrastructure as a service (IaaS) model, the cloud consumer has responsibility and control over the operating system.

In the software as a service (SaaS) model, the cloud service provider is responsible for operating systems.

Public and private cloud are cloud deployment models, not cloud service models.

96.

Which of the following can be added to the PCI slot of a SAN disk array or switch to offload processing load and improve performance?

HBA

Router

NAS

RAID

Correct answer: HBA

An HBA (host bus adapter) is typically installed as a PCO add-on card. HBAs can increase performance by offloading processing load needed for storage data. As a result, the storage system can allocate its processing power to operating system-level functions instead of storage input/output.

A router is a type of network device that transmits layer 3 (network) traffic.

A network attached storage (NAS) is another type of storage system.

Redundant array of independent disks (RAID) is a storage virtualization approach that can improve data resilience.

97.

You need to implement tunneling for Ethernet frames over an IP network. What tunneling protocol should you use?

L2TP

IKEv1

TCP SYN

IKEv2

Correct answer: L2TP

L2TP is a tunneling protocol that supports Ethernet frames and a variety of other packet types. It does not support encryption natively but can be used with IPsec if encryption is needed.

IKE (internet key exchange) v1 and v2 are used for negotiation and authentication with IPsec but are not tunneling protocols themselves.

TCP SYN is part of the 3-way TCP handshake, not a tunneling protocol.

98.

Acme Inc. management wants to understand the usage of each of their private cloud services so they can make budgeting decisions.

What can Acme Inc. use to achieve this goal?

Showbacks

SLAs

MTTR

MTBF

Correct answer: Showbacks

Showbacks track usage of IT services to enable decision-making about value relative to cost. Services that are rarely used may not be worth further investment.

A service level agreement (SLA) is a contractual agreement to meet a specific standard for a service. For example, SLAs often define availability and uptime requirements.

Mean time between failures (MTBF) and mean time to repair (MTTR) are common system health and reliability metrics.

99.

Which of the following is NOT a relational database?

MongoDB

MySQL

Oracle Database

IBM DB2

Correct answer: MongoDB

MongoDB is a nonrelational database.

MySQL, Oracle Database, and IBM DB2 are all relational databases.

100.

Which of the following statements about dedicated versus shared compute environments is FALSE?

Dedicated compute environments are typically cheaper

Dedicated compute environments typically offer more consistent performance

Dedicated compute environments may be required due to regulatory requirements

Shared compute environments allow cloud service providers to distribute costs across multiple consumers

Correct answer: Dedicated compute environments are typically cheaper

Dedicated compute environments are typically more expensive than shared compute environments because the costs cannot be distributed across multiple consumers.

Dedicated compute environments typically offer more consistent performance than shared compute environments.

Dedicated compute environments may be required due to regulatory requirements.

Shared compute environments allow cloud service providers to distribute costs across multiple consumers.
