DUMPSARENA

Certificate of Cloud Security Knowledge (v5.0)

Cloud Security Alliance CCSK

Version Demo

Total Demo Questions: 35

Total Premium Questions: 355

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Computing Concepts and Architectures	58
Topic 2, Cloud Security Governance and Risk Management	42
Topic 3, Legal Issues, Contracts and Electronic Discovery	12
Topic 4, Compliance and Audit Management	22
Topic 5, Information Governance	4
Topic 6, Management Plane and Business Continuity	51
Topic 7, Infrastructure Security	29
Topic 8, Virtualization and Container Security	26
Topic 9, Application Security	26
Topic 10, Data Security and Encryption	27
Topic 11, Security as a Service	10
Topic 12, Identity, Entitlement, and Access Management	48
Total	355

QUESTION NO: 1

Which of the following is NOT a cloud computing characteristic that impacts incidence response?

A. The on demand self-service nature of cloud computing environments.

B. Privacy concerns for co-tenants regarding the collection and analysis of telemetry and artifacts associated with an incident.

C. The possibility of data crossing geographic or jurisdictional boundaries.

D. Object-based storage in a private cloud.

E. The resource pooling practiced by cloud services, in addition to the rapid elasticity offered by cloud infrastructures.

ANSWER: D

Explanation:

Object-based storage in a private cloud. is the correct selection because it is a particular storage design and deployment implementation, rather than a fundamental characteristic of cloud computing. Object storage organizes data as objects with associated metadata and is commonly used by cloud providers, but its presence is neither required for a service to be considered cloud computing nor uniquely relevant to incident-response operations. An organization can operate a private cloud using block, file, object, or other storage architectures. Incident responders may need to understand whichever storage technology is in use, including its logging, retention, access controls, and evidence-collection capabilities; however, that is an environment-specific technical consideration rather than a cloud characteristic. Cloud-computing characteristics are generally defined in terms of service behavior and operational properties, such as self-service provisioning, shared resources, elasticity, and broad network access. These properties can affect evidence availability, ownership boundaries, speed of environmental change, and coordination during an incident. The NIST cloud-computing definition distinguishes these essential characteristics from implementation technologies and deployment details. See the NIST Definition of Cloud Computing (SP 800-145) and the Cloud Security Alliance's Security Guidance for Critical Areas of Focus in Cloud Computing.

QUESTION NO: 2

What are the key outcomes of implementing robust cloud risk management practices?

A. Ensuring the security and resilience of cloud environments

B. Negotiating shared responsibilities

C. Transferring compliance to the cloud service provider via inheritance

D. Reducing the need for compliance with regulatory requirements

ANSWER: A

Explanation:

Ensuring the security and resilience of cloud environments is the key outcome of robust cloud risk management. Cloud risk management is a continuous organizational process for identifying relevant threats, vulnerabilities, and business impacts; assessing the likelihood and magnitude of risk; selecting and implementing appropriate treatments; and monitoring whether controls remain effective as the cloud environment, service configuration, and threat landscape change. Applied effectively, it helps an organization protect the confidentiality, integrity, and availability of its cloud-hosted information and services while maintaining the ability to withstand, respond to, and recover from disruptive events. In cloud environments, this includes understanding the provider's and customer's respective control responsibilities, validating that required safeguards are implemented, and maintaining governance over data, identities, configurations, third parties, and compliance obligations. The Cloud Security Alliance Cloud Controls Matrix supports this outcome by providing a control framework organizations can use to assess cloud-security risk and manage assurance across cloud services. NIST's risk-management guidance likewise frames risk management as an ongoing activity that supports trustworthy, resilient systems and informed business decisions. See the Cloud Security Alliance Cloud Controls Matrix and NIST Special Publication 800-39.

QUESTION NO: 3

Which technique involves assessing potential threats through analyzing attacker capabilities, motivations, and potential targets?

A. Threat modeling

B. Vulnerability assessment

C. Incident response

D. Risk assessment

ANSWER: A

Explanation:

Threat modeling is the structured, proactive technique for identifying and evaluating how a system could be attacked by considering relevant threat actors, their capabilities, their motivations, likely attack paths, and the assets or components they may target. It places the application, workload, or business process in context: teams identify what must be protected, map trust boundaries and data flows, consider plausible adversarial actions, and prioritize mitigations according to the threats that matter most. This makes threat modeling useful early in architecture and design, as well as when significant changes are made to a cloud environment. In cloud security, the process helps ensure that controls are selected based on realistic threats to cloud data, identities, interfaces, workloads, and shared-responsibility boundaries rather than only on generic checklists. The Cloud Security Alliance describes threat modeling as a key practice for secure application design and risk management, and emphasizes incorporating security throughout the development lifecycle. Additional practical guidance on conducting threat modeling, including identifying threats and defining mitigations, is available from OWASP Threat Modeling. For cloud-specific security governance and control context, see the Cloud Security Alliance Cloud Controls Matrix.

QUESTION NO: 4

Containers are highly portable code execution environments.

A. False

B. True

ANSWER: B

Explanation:

True is correct. A container packages an application together with the runtime components it needs, such as libraries, dependencies, configuration, and filesystem content, into a standardized image. That image can be run consistently by a compatible container runtime across different environments, including a developer workstation, private cloud, or public-cloud infrastructure. This packaging model reduces the differences between deployment environments that historically caused applications to behave differently after being moved. Portability is therefore a core operational characteristic of containers and is one reason they are widely used for modern application delivery and orchestration platforms. The portability applies to the packaged application environment rather than eliminating all platform dependencies: the target must provide a compatible operating-system kernel and container runtime. From a cloud-security perspective, this portability also means that image provenance, secure image registries, vulnerability management, and configuration controls should travel with and be assessed throughout the container lifecycle. NIST describes containers as a form of operating-system-level virtualization that packages applications and their dependencies, while CSA guidance addresses the security considerations for these portable cloud-native workloads. See NIST SP 800-190, Application Container Security Guide and the Cloud Security Alliance Security Guidance v4.

QUESTION NO: 5

In which deployment model should the governance strategy consider the minimum common set of controls comprised of the Cloud Service Provider contract and the organization's internal governance agreements?

A. Public

B. PaaS

C. Private

D. IaaS

E. Hybrid

ANSWER: E

Explanation:

Hybrid is correct because a hybrid cloud combines distinct cloud environments, commonly involving a mix of private infrastructure and services operated by one or more cloud service providers. Governance must therefore work across organizational boundaries and differing operating models. The effective baseline is the minimum common set of controls that can be consistently enforced and evidenced across those environments. That baseline is shaped both by the controls and commitments formally established in the cloud service provider contract and by the organization's own internal governance requirements, such as risk appetite, policy obligations, data-classification rules, and compliance responsibilities.

This common-control approach enables the organization to maintain coherent oversight even where technical implementation, ownership, and responsibility vary between environments. It also supports clear accountability under the shared-responsibility model: contractual terms define provider commitments, while internal agreements establish how the organization governs its people, processes, data, and use of cloud services. CSA's guidance emphasizes that cloud governance must align contractual, operational, and assurance requirements to manage risk across cloud arrangements. See the Cloud Security Alliance Security Guidance v4 and the Cloud Controls Matrix.

QUESTION NO: 6

Which factors primarily drive organizations to adopt cloud computing solutions?

A. Scalability and redundancy

B. Improved software development methodologies

C. Enhanced security and compliance

D. Cost efficiency and speed to market

ANSWER: D

Explanation:

Cost efficiency and speed to market are primary business drivers for cloud adoption because cloud services shift organizations away from large up-front capital expenditures and toward consumption-based operating costs. This model can reduce the cost and time associated with acquiring, deploying, maintaining, and refreshing infrastructure. It also lets teams provision resources rapidly rather than waiting for traditional procurement and data-center deployment cycles. As a result, organizations can experiment, build, release, and scale products or services more quickly, directly improving business agility and time to market.

Cloud Security Alliance guidance identifies economic benefits and increased agility as central cloud value propositions. The ability to obtain infrastructure and platform capabilities on demand enables organizations to respond to changing customer needs and market conditions with greater speed. These benefits are not limited to technology modernization; they are strategic outcomes that commonly motivate the decision to adopt cloud computing. See the Cloud Security Alliance Security Guidance v4.0 and its cloud security guidance resources.

QUESTION NO: 7

Which of the following is a common risk factor related to misconfiguration and inadequate change control in cybersecurity?

A. Failure to update access controls after employee role changes

B. Lack of sensitive data encryption

C. Lack of 3rd party service provider specialized in patch management procedures

D. Excessive SBOM focus

ANSWER: A

Explanation:

Failure to update access controls after employee role changes is a common and material risk arising from weak change control and configuration management. When a person transfers teams, changes responsibilities, or leaves an organization, their permissions must be reviewed and adjusted promptly to match their current business need. If this change is not controlled through a defined joiner-mover-leaver process, formerly appropriate privileges can remain active. Those stale permissions can enable unauthorized access to cloud resources, sensitive information, administrative functions, or systems outside the person's current role.

In cloud environments, access configurations are implemented through identities, groups, roles, policies, and entitlements, often across multiple services. Effective change management therefore requires authorized requests, approval, implementation, validation, and periodic access reviews. This supports least privilege and ensures that access-control settings remain aligned with organizational changes rather than accumulating over time. The Cloud Security Alliance Cloud Controls Matrix identifies identity and access management controls addressing authorization, least privilege, and timely revocation or modification of access. Guidance on access-control lifecycle practices is also provided by Cloud Security Alliance Cloud Controls Matrix and NIST SP 800-53 Rev. 5.

QUESTION NO: 8

In the shared security model, how does the allocation of responsibility vary by service?

A. Shared responsibilities should be consistent across all services.

B. Based on the per-service SLAs for security.

C. Responsibilities are the same across IaaS, PaaS, and SaaS in the shared model.

D. Responsibilities are divided between the cloud provider and the customer based on the service type.

ANSWER: D

Explanation:

Responsibilities are divided between the cloud provider and the customer based on the service type. This is the core principle of the cloud shared-responsibility model: the provider is responsible for securing the underlying cloud infrastructure, while the customer retains responsibility for security *in* the cloud, including areas under its control. The precise boundary shifts with the service model. In IaaS, customers generally manage and secure guest operating systems, applications, identities, configurations, and their data. With PaaS, the provider assumes more responsibility for the platform components, while the customer focuses on applications, data, access management, and secure use of platform services. In SaaS, the provider operates most of the application stack, but customers remain accountable for such matters as data governance, user access, identity configuration, and appropriate configuration of the service. The allocation must therefore be understood for each service and documented through the provider's terms, security documentation, and the customer's own governance processes. It is not a fixed division that applies identically to every cloud offering. The Cloud Security Alliance describes this model as a division of responsibilities that depends on the cloud service delivery model and the services consumed. See the Cloud Security Alliance Security Guidance v4 and the Cloud Controls Matrix.

QUESTION NO: 9

CCM: A hypothetical company called: "Health4Sure" is located in the United States and provides cloud based services for tracking patient health. The company is compliant with HIPAA/HITECH Act among other industry standards. Health4Sure

decides to assess the overall security of their cloud service against the CCM toolkit so that they will be able to present this document to potential clients.

Which of the following approach would be most suitable to assess the overall security posture of Health4Sure's cloud service?

A. The CCM columns are mapped to HIPAA/HITECH Act and therefore Health4Sure could verify the CCM controls already covered ad a result of their compliance with HIPPA/HITECH Act. They could then assess the remaining controls. This approach will save time.

B. The CCM domain controls are mapped to HIPAA/HITECH Act and therefore Health4Sure could verify the CCM controls already covered as a result of their compliance with HIPPA/HITECH Act. They could then assess the remaining controls thoroughly. This approach saves time while being able to assess the company's overall security posture in an efficient manner.

C. The CCM domains are not mapped to HIPAA/HITECH Act. Therefore Health4Sure should assess the security posture of their cloud service against each and every control in the CCM. This approach will allow a thorough assessment of the security posture.

ANSWER: B

Explanation:

The CCM domain controls are mapped to HIPAA/HITECH Act and therefore Health4Sure could verify the CCM controls already covered as a result of their compliance with HIPAA/HITECH Act. They could then assess the remaining controls thoroughly. This approach saves time while being able to assess the company's overall security posture in an efficient manner. The Cloud Controls Matrix is designed to support this type of cross-reference: its controls are mapped to numerous laws, regulations, standards, and frameworks, including HIPAA-related requirements. Health4Sure can use its existing HIPAA/HITECH compliance evidence as a starting point to identify CCM requirements for which relevant policies, procedures, technical safeguards, assessments, and audit artifacts may already exist. The organization must still validate that the evidence satisfies the specific CCM control intent and perform a complete assessment of CCM controls not adequately covered by the HIPAA/HITECH mapping. This produces a defensible CCM-aligned security assessment without unnecessarily duplicating work already performed for its regulated healthcare compliance program. CCM mappings are an efficiency and scoping mechanism; the resulting assessment should document the evidence and the status of every applicable CCM control for prospective customers. See the CSA's Cloud Controls Matrix v4 and the HHS overview of the HIPAA Security Rule.

QUESTION NO: 10

Which security capability provides a record of cloud management actions, such as creating a storage repository or changing an identity policy?

A. Management-plane audit logging

B. Data fragmentation

C. Client-side encryption

D. Resource pooling

E. Endpoint disk sanitization

ANSWER: A

Explanation:

Management-plane audit logging is correct because it records administrative and API actions performed through the cloud provider's management interfaces. These records can identify the initiating identity, the action requested, the target resource, the time of the event, the source location, and whether the action succeeded.

Management-plane audit logs are essential for detecting unauthorized configuration changes, investigating account compromise, demonstrating accountability, and supporting incident response. Logging must be enabled, protected against

unauthorized modification, retained for an appropriate period, and monitored for suspicious activity. It is distinct from application logging, which records events generated by the customer application itself. See the Cloud Controls Matrix v4 and NIST SP 800-92.

QUESTION NO: 11

A company plans to shift its data processing tasks to the cloud. Which type of cloud workload best describes the use of software emulations of physical computers?

A. Platform as a Service (PaaS)

B. Serverless Functions (FaaS)

C. Containers

D. Virtual Machines (VMs)

ANSWER: D

Explanation:

Virtual Machines (VMs) is correct because a virtual machine is a software-defined computer that emulates the essential hardware resources of a physical system, including compute, memory, storage, and networking. A hypervisor virtualizes the underlying physical server and permits separate guest operating systems and their applications to execute in isolated VM environments. This model is particularly appropriate when a company's data-processing workload needs an entire operating system, requires specific machine-level configuration, or must retain compatibility with existing applications designed for conventional servers.

In cloud environments, VMs are typically delivered as infrastructure-as-a-service resources. The cloud provider is responsible for operating and securing the underlying physical infrastructure and virtualization layer, while the customer commonly retains responsibility for guest operating system configuration, workload security, identity controls, application patching, and data protection. This division is important for cloud security planning because migrating a server workload to a VM does not eliminate the need to harden and manage its operating system. The CSA describes IaaS as providing fundamental computing resources on which customers deploy and run software, including operating systems and applications. NIST also defines virtualization as an abstraction of physical computing resources that enables virtual machines. See the Cloud Security Alliance Security Guidance v4 and NIST SP 800-145.

QUESTION NO: 12

What is a commonly used method by which hybrid cloud integrates data centers with public cloud?

A. Using VPN or dedicated links

B. Using peer-to-peer networks

C. Using local area network (LAN)

D. Using satellite connections

ANSWER: A

Explanation:

Using VPN or dedicated links is a standard method for integrating an organization's on-premises data center with public-cloud resources in a hybrid-cloud architecture. A site-to-site VPN creates an encrypted tunnel over the public internet, allowing private network traffic to move securely between the data center and a cloud virtual network. Dedicated connectivity provides a private, provider-supported path between the organization and the cloud provider, commonly offering more predictable bandwidth, latency, availability, and performance characteristics for sustained or sensitive workloads. These connections enable systems in the two environments to communicate as part of an integrated architecture while retaining workloads, data, or services in the locations that best meet business, security, operational, and regulatory requirements. The

Cloud Security Alliance describes hybrid deployment as a composition of two or more distinct cloud infrastructures that remain unique entities but are bound together by technology enabling data and application portability. Secure network connectivity is therefore a foundational enabler of hybrid integration. See the CSA Security Guidance v4 and the NIST definition of hybrid cloud in SP 800-145.

QUESTION NO: 13

The Software Defined Perimeter (SDP) includes which components?

A. Client, Controller, and Gateway

B. Client, Controller, Firewall, and Gateway

C. Client, Firewall, and Gateway

D. Controller, Firewall, and Gateway

E. Client, Controller, and Firewall

ANSWER: A

Explanation:

Client, Controller, and Gateway is correct because these are the core logical components defined by the Cloud Security Alliance SDP architecture. The SDP client runs on, or represents, the requesting device or workload and initiates authenticated connection requests. The SDP controller is the policy decision and control-plane component: it authenticates users and devices, evaluates authorization policy and context, and determines whether a connection to a protected resource may be established. The SDP gateway, sometimes called an accepting host, enforces that decision at the protected-service side. It receives controller instructions and permits only authorized, mutually authenticated, encrypted connections to the application or infrastructure it protects.

This separation enables an identity- and context-driven, need-to-know access model. Protected services are not broadly discoverable or reachable before the controller authorizes access, which reduces exposed attack surface and supports zero-trust principles. A firewall may coexist with an SDP deployment as a complementary network-security control, but it is not one of the three fundamental SDP architectural components. CSA describes this client, controller, and gateway model in its Software Defined Perimeter Specification and relates SDP to zero-trust architecture in its SDP and Zero Trust Network Architecture guidance.

QUESTION NO: 14

What is the primary function of a Load Balancer Service in a Software Defined Network (SDN) environment?

A. To create isolated virtual networks

B. To monitor network performance and activity

C. To distribute incoming network traffic across multiple destinations

D. To encrypt data for secure transmission

ANSWER: C

Explanation:

To distribute incoming network traffic across multiple destinations is correct because load balancing provides a virtual service endpoint that receives client requests and selects an appropriate backend destination, such as an application server, virtual machine, container, or service instance. In an SDN environment, the control plane can program network behavior centrally, allowing traffic-distribution policies to be applied dynamically in response to topology, workload, service health, and demand. The objective is to avoid concentrating requests on a single destination while improving service availability, capacity utilization, and application responsiveness. A load balancer commonly uses a defined algorithm, such as round robin or least connections, together with health checks to send new traffic only to destinations able to serve it. Load-

balancing services may operate at the transport or application layer and can also support features such as session persistence and TLS termination, but their core purpose remains selecting and distributing flows among multiple available endpoints. This aligns with the Cloud Security Alliance guidance that cloud network services are software-defined and policy-driven, and that resilient cloud architectures use distributed resources and automated control capabilities. See the Cloud Security Alliance Security Guidance v4 and the AWS Elastic Load Balancing documentation.

QUESTION NO: 15

What is true of companies considering a cloud computing business relationship?

A. The laws protecting customer data are based on the cloud provider and customer location only.

B. The confidentiality agreements between companies using cloud computing services is limited legally to the company, not the provider.

C. The companies using the cloud providers are the custodians of the data entrusted to them.

D. The cloud computing companies are absolved of all data security and associated risks through contracts and data laws.

E. The cloud computing companies own all customer data.

ANSWER: C

Explanation:

The companies using the cloud providers are the custodians of the data entrusted to them. An organization that collects, creates, or otherwise controls information remains accountable to the individuals, business partners, regulators, and contractual counterparties that entrusted that information to it, even after processing or storage is outsourced to a cloud service provider. A cloud arrangement changes the operational model and allocates responsibilities between the customer and provider; it does not remove the customer organization's overarching governance, privacy, security, and compliance obligations. Consequently, before entering a cloud business relationship, the customer should establish appropriate contractual commitments, assess the provider's security controls, define responsibilities under the shared-responsibility model, and maintain oversight of how entrusted data is handled throughout its lifecycle. The Cloud Security Alliance Cloud Controls Matrix is designed to support this due diligence by mapping cloud control responsibilities across domains such as data security, governance, and legal requirements. CSA guidance also emphasizes that cloud customers need effective governance and assurance processes for provider-managed services. See the Cloud Controls Matrix and the CSA Security Guidance.

QUESTION NO: 16

What is a potential concern of using Security-as-a-Service (SecaaS)?

A. Lack of visibility

B. Deployment flexibility

C. Scaling and costs

D. Intelligence sharing

E. Insulation of clients

ANSWER: A

Explanation:

Lack of visibility is a potential concern when using Security-as-a-Service (SecaaS) because important security capabilities, telemetry, and operational processes are delivered and operated through a provider-managed service. The customer may not have direct access to the underlying infrastructure, complete event data, detection logic, administrative activity, configuration details, or evidence of how the provider performs specific controls. This can make it harder to maintain

continuous situational awareness, validate that security requirements are being met, investigate incidents, and provide audit evidence to regulators or internal assurance teams.

Visibility must therefore be explicitly addressed during provider selection and contract design. Customers should define required logging, reporting, alerting, API access, data-retention periods, incident-notification procedures, audit rights, control attestations, and clear shared-responsibility boundaries. These requirements allow the customer to retain appropriate governance even though the service provider operates much of the security function. CSA guidance emphasizes that cloud customers need transparency into provider controls and sufficient assurance mechanisms to manage cloud risks effectively. See the Cloud Security Alliance Security Guidance v4 and CSA's cloud security guidance resources.

QUESTION NO: 17

In a containerized environment, what is fundamental to ensuring runtime protection for deployed containers?

A. Implementing real-time visibility

B. Deploying container-specific antivirus scanning

C. Using static code analysis tools in the pipeline

D. Full packet network monitoring

ANSWER: A

Explanation:

Implementing real-time visibility is fundamental because container runtime protection depends on continuously observing what deployed workloads actually do after they start. Containers are dynamic, short-lived, and commonly orchestrated at scale, so point-in-time checks performed before deployment cannot by themselves identify activity that emerges in production. Runtime visibility supplies telemetry on executing processes, container images and identities, system calls, network connections, configuration changes, and interactions with the orchestration environment. This context enables security teams and automated controls to detect anomalous or policy-violating behavior promptly, investigate the affected workload, and take a proportionate response such as alerting, isolating a container, or terminating a malicious process. It also supports enforcement of runtime policies based on the expected behavior of a particular workload rather than relying solely on broad infrastructure assumptions. The Cloud Security Alliance emphasizes workload-security capabilities that provide visibility and monitoring throughout the workload lifecycle, including runtime operation. In a containerized environment, this continuous operational insight is therefore the basis for detecting and responding to threats against deployed containers. See the Cloud Security Alliance Security Guidance v5 and the CSA's Definitive Guide to Cloud Security for cloud workload security guidance.

QUESTION NO: 18

CCM: The following list of controls belong to which domain of the CCM?

GRM 06 – Policy GRM 07 – Policy Enforcement GRM 08 – Policy Impact on Risk Assessments GRM 09 – Policy Reviews GRM 10 – Risk Assessments GRM 11 – Risk Management Framework

A. Governance and Retention Management

B. Governance and Risk Management

C. Governing and Risk Metrics

ANSWER: B

Explanation:

Governance and Risk Management is correct because the control identifiers use the **GRM** prefix, which denotes the Cloud Controls Matrix domain addressing governance and risk-management practices. The listed controls cover the lifecycle of organizational policy and risk oversight: establishing policy, enforcing it, assessing its impact on risk assessments, reviewing

it on an ongoing basis, conducting risk assessments, and maintaining a risk-management framework. These activities provide the governance structure through which a cloud service provider defines accountability, makes risk-informed decisions, and ensures that security requirements are managed consistently over time. In the CCM domain model used by this question, GRM is explicitly the abbreviation for Governance and Risk Management. This domain is intended to ensure that security governance is documented, operationalized, reviewed, and integrated with enterprise risk-management processes rather than handled as an isolated technical activity. The Cloud Security Alliance publishes the CCM as a controls framework for assessing cloud security and for mapping cloud controls to widely used standards and regulatory frameworks. See the Cloud Security Alliance Cloud Controls Matrix overview and the CSA's CCM v3.0.1 artifact.

QUESTION NO: 19

Which of the following best describes an aspect of PaaS services in relation to network security controls within a cloud environment?

A. They override the VNet/VPC's network security controls by default

B. They do not interact with the VNet/VPC's network security controls

C. They require manual configuration of network security controls, separate from the VNet/VPC

D. They often inherit the network security controls of the underlying VNet/VPC

ANSWER: D

Explanation:

They often inherit the network security controls of the underlying VNet/VPC is the best description because PaaS is delivered on provider-managed infrastructure while still commonly being connected to, exposed through, or privately integrated with a customer-controlled virtual network. The customer can therefore apply network policy at shared enforcement points, such as subnet boundaries, routing paths, private endpoints, security groups, network security groups, firewall rules, and access-control mechanisms. This supports a consistent defense-in-depth model: application platform capabilities remain managed by the cloud provider, while the customer retains responsibility for configuring the network exposure and permitted communication paths within its cloud tenancy.

The precise implementation depends on the provider and PaaS product. For example, a service may be reachable through a private endpoint, use VNet/VPC integration for egress, or be governed by service-specific firewall settings in addition to network policy. Nevertheless, the important PaaS security principle is that platform abstraction does not eliminate the relevance of virtual-network controls. Properly designed PaaS deployments can use those controls to restrict traffic to approved sources, destinations, and paths. This aligns with the Cloud Security Alliance's emphasis on clearly assigning shared-responsibility controls and protecting cloud network boundaries. See the Cloud Security Alliance Security Guidance v4 and Azure App Service VNet integration documentation.

QUESTION NO: 20

In federated identity management, what role does the identity provider (IdP) play in relation to the relying party?

A. The IdP relies on the relying party to authenticate and authorize users.

B. The relying party makes assertions to the IdP about user authorizations.

C. The IdP and relying party have no direct trust relationship.

D. The IdP makes assertions to the relying party after building a trust relationship.

ANSWER: D

Explanation:

The IdP makes assertions to the relying party after building a trust relationship. In a federated identity model, the identity provider is the authoritative party that authenticates the user and produces a signed assertion or token containing identity

information and, where applicable, authentication context and attributes. The relying party, commonly an application or service provider, is configured to trust assertions from that IdP. This trust is established in advance through federation metadata, signing certificates or keys, registered endpoints, issuer and audience values, and agreed protocol settings.

After authenticating the user, the IdP sends an assertion to the relying party using a federation protocol such as SAML or OpenID Connect. The relying party validates the assertion's signature and relevant claims before using it to establish an application session and make access-control decisions. This arrangement supports single sign-on while allowing the application to avoid directly handling the user's primary credentials. The Cloud Security Alliance describes identity federation as a mechanism that enables authentication information to be exchanged across distinct security domains, with trust and assertion validation central to the model. See the Cloud Security Alliance Security Guidance v4 and the OASIS SAML 2.0 Core specification.

QUESTION NO: 21

Use elastic servers when possible and move workloads to new instances.

A. False

B. True

ANSWER: B

Explanation:

"True" is correct because cloud elasticity supports an immutable, replace-rather-than-repair operating model. Workloads can be deployed to newly provisioned instances created from an approved, hardened image or infrastructure-as-code definition, rather than relying on long-lived servers that accumulate untracked changes. This reduces configuration drift and helps ensure that security patches, baseline settings, endpoint controls, logging configuration, and application releases are consistently incorporated into each new instance.

Moving workloads to replacement instances also limits the lifetime of a potentially compromised or misconfigured host. Once traffic and state have been safely transitioned, the previous instance can be terminated, removing its active credentials, processes, and storage exposure according to the organization's retention and sanitization requirements. Elastic capacity makes this practical because replacement capacity can be created quickly and scaled to maintain availability during maintenance, remediation, or deployment. This practice aligns with cloud security guidance emphasizing automation, repeatable secure baselines, and use of cloud-native capabilities to improve operational resilience. See the Cloud Security Alliance's Security Guidance v4 and NIST's Application Container Security Guide for related guidance on automated, repeatable deployment and lifecycle management.

QUESTION NO: 22

Which approach is essential in identifying compromised identities in cloud environments where attackers utilize automated methods?

A. Focusing exclusively on signature-based detection for known malware

B. Deploying behavioral detectors for IAM and management plane activities

C. Implementing full packet capture and monitoring

D. Relying on IP address and connection header monitoring

ANSWER: B

Explanation:

Deploying behavioral detectors for IAM and management plane activities is essential because identity-focused cloud attacks often use valid, compromised credentials and automated tooling rather than malware that can be recognized by a static signature. Effective detection therefore depends on establishing a baseline for normal identity behavior and alerting on meaningful deviations: unusual API-call sequences, impossible travel, anomalous login patterns, privilege-escalation

attempts, creation of new access keys or roles, access from unfamiliar locations or devices, and high-volume or atypical management-plane activity. Cloud IAM and control-plane audit logs provide the telemetry needed to correlate these events across identities, sessions, permissions, and resources. Behavioral analytics can also identify attack automation through patterns such as rapid enumeration, repeated failed operations, or systematic use of cloud APIs. This approach supports continuous monitoring and incident response while remaining effective when an attacker's actions are performed through otherwise legitimate cloud interfaces. The Cloud Security Alliance emphasizes logging, monitoring, and detection capabilities as foundational cloud-security controls, including visibility into administrative and identity activity. Guidance for AWS similarly highlights monitoring management events through CloudTrail, which records actions taken in an account and supplies core evidence for detecting suspicious IAM and control-plane behavior. See the Cloud Controls Matrix and AWS CloudTrail User Guide.

QUESTION NO: 23

Select the best definition of "compliance" from the options below.

A. The development of a routine that covers all necessary security measures.

B. The diligent habits of good security practices and recording of the same.

C. The timely and efficient filing of security reports.

D. The awareness and adherence to obligations, including the assessment and prioritization of corrective actions deemed necessary and appropriate.

E. The process of completing all forms and paperwork necessary to develop a defensible paper trail.

ANSWER: D

Explanation:

The awareness and adherence to obligations, including the assessment and prioritization of corrective actions deemed necessary and appropriate, is the best definition of compliance because compliance is an ongoing governance activity rather than a one-time paperwork or reporting exercise. Organizations must identify the legal, regulatory, contractual, and internal-policy obligations that apply to their cloud services; understand how those obligations affect their controls and operations; and demonstrate that required safeguards are being maintained. Where assessments identify gaps, compliance also requires risk-informed remediation: corrective actions should be evaluated, prioritized based on impact and urgency, assigned to accountable owners, and tracked through completion. This approach makes compliance part of continuous control monitoring and improvement, supporting evidence-based assurance to customers, auditors, and regulators. Cloud Security Alliance guidance emphasizes that cloud governance must address applicable laws, regulations, standards, contractual commitments, and the controls used to meet them. The Cloud Controls Matrix likewise provides a control framework organizations can use to assess and manage cloud-security compliance obligations. See the Cloud Security Alliance Security Guidance and the Cloud Controls Matrix.

QUESTION NO: 24

In volume storage, what method is often used to support resiliency and security?

A. proxy encryption

B. data rights management

C. hypervisor agents

D. data dispersion

E. random placement

ANSWER: D

Explanation:

Data dispersion is correct because it protects volume-storage data by splitting it into fragments and distributing those fragments across separate storage locations or nodes. A dispersion design can apply redundancy or erasure coding so that the original volume remains recoverable when a storage node, disk, or location becomes unavailable. This directly contributes to resiliency by reducing dependence on any single storage component and enabling continued access or reconstruction after failures.

It also supports security because an individual fragment does not ordinarily provide a usable copy of the underlying data. Reconstructing the volume requires the necessary collection of fragments and the applicable reconstruction process, which reduces the impact of exposure involving one storage location. In cloud environments, this approach can help address both availability and confidentiality considerations for distributed storage, provided the organization also manages access controls, encryption keys, and recovery thresholds appropriately. The Cloud Security Alliance identifies data storage, availability, and protection mechanisms as core considerations in cloud security architecture. See the Cloud Security Alliance Security Guidance for Critical Areas of Focus in Cloud Computing and the Cloud Controls Matrix.

QUESTION NO: 25

What can be implemented to help with account granularity and limit

blast radius with IaaS an PaaS?

A. Configuring secondary authentication

B. Establishing multiple accounts

C. Maintaining tight control of the primary account holder credentials

D. Implementing least privilege accounts

E. Configuring role-based authentication

ANSWER: B

Explanation:

Establishing multiple accounts is correct because an account boundary is a strong administrative and security isolation mechanism in IaaS and PaaS environments. Separating workloads, environments, business units, or risk domains into distinct accounts creates independent scopes for identity policies, resource ownership, logging, billing, quotas, and service configuration. Consequently, an error, compromised credential, overly broad permission, or misconfigured resource in one account is less likely to expose or disrupt resources assigned to another account. This reduces the potential blast radius while also providing the account-level granularity needed to delegate administration and apply controls appropriate to each workload. The Cloud Security Alliance emphasizes defining logical isolation and appropriate identity-and-access-management boundaries as part of cloud governance. Major cloud-provider security guidance similarly treats a multi-account structure as a foundational way to isolate workloads and contain security incidents. See the Cloud Security Alliance Cloud Controls Matrix and AWS guidance on reducing permissions and isolating workloads through account boundaries.

QUESTION NO: 26

Which cloud-based service model enables companies to provide client-based access for partners to databases or applications?

A. Platform-as-a-service (PaaS)

B. Desktop-as-a-service (DaaS)

C. Infrastructure-as-a-service (IaaS)

D. Identity-as-a-service (IDaaS)

E. Software-as-a-service (SaaS)

ANSWER: A

Explanation:

Platform-as-a-service (PaaS) is correct because it supplies a managed application platform on which an organization can build, deploy, integrate, and operate applications for its users, customers, and business partners. The provider manages the underlying cloud infrastructure and much of the platform stack, while the customer retains control of the applications it creates or deploys. This lets a company expose a purpose-built application or integration endpoint to partners through a client interface, while using managed platform capabilities such as application runtimes, middleware, data services, APIs, and database integration. In this context, “client-based access” means that partners can use an application client or browser-based interface to interact with functionality and data delivered by the organization’s deployed solution; it does not require partners to administer the hosting environment. The NIST cloud-computing definition describes PaaS as the capability for a consumer to deploy consumer-created or acquired applications onto the provider’s cloud infrastructure. CSA’s security guidance likewise treats PaaS as a model that provides the platform components needed to develop, deploy, and run cloud applications. See the NIST SP 800-145 cloud service-model definitions and the Cloud Security Alliance Security Guidance v4.

QUESTION NO: 27

Which of the following cloud essential characteristics refers to the capability of the service to scale resources up or down quickly and efficiently based on demand?

A. On-Demand Self-Service

B. Broad Network Access

C. Resource Pooling

D. Rapid Elasticity

ANSWER: D

Explanation:

Rapid Elasticity is the cloud essential characteristic that describes the ability to provision and release computing capabilities quickly, sometimes automatically, in response to changing demand. This permits a cloud customer’s available compute, storage, or network capacity to expand when workload demand rises and contract when demand falls. To the consumer, the resources may appear effectively unlimited and can be acquired in quantities appropriate to the current need. This dynamic scaling capability is central to cloud’s operational flexibility: rather than maintaining infrastructure sized for peak demand at all times, an organization can use resources as demand requires and reduce allocated capacity afterward. NIST’s widely adopted cloud-computing definition identifies rapid elasticity as one of the five essential characteristics and specifically notes that capabilities can be elastically provisioned and released to scale rapidly outward and inward commensurate with demand. CSA CCSK materials use these established cloud characteristics as foundational terminology for understanding cloud services, architectures, and their associated governance and security considerations. See the NIST SP 800-145 definition of cloud computing and the Cloud Security Alliance Security Guidance v4.

QUESTION NO: 28

What does it mean if the system or environment is built automatically from a template?

A. Nothing.

B. It depends on how the automation is configured.

C. Changes made in production are overwritten by the next code or template change.

D. Changes made in test are overwritten by the next code or template change.

E. Changes made in production are untouched by the next code or template change.

ANSWER: C

Explanation:

“Changes made in production are overwritten by the next code or template change.” is correct because an environment built automatically from a template follows an infrastructure-as-code and configuration-management model. The template and its associated code represent the intended, authoritative configuration of the environment. When the automated deployment process runs again—whether to update the application, apply a revised template, scale resources, or rebuild an instance—it recreates or reconciles the environment with that declared state. Consequently, manual production changes that were not incorporated into the template or deployment code are not durable and can be replaced during the next deployment.

This approach supports repeatability, traceability, controlled change management, and consistent security configuration. Operational changes should therefore be made through the approved source-controlled template and automation pipeline, then promoted and deployed, rather than performed only as ad hoc production modifications. This is a core benefit of automated provisioning: environments can be consistently rebuilt from known definitions. CSA guidance emphasizes the importance of controlled, repeatable configuration and change processes in cloud environments, while NIST configuration-management guidance similarly treats authorized configuration baselines as the basis for maintaining systems. See the Cloud Security Alliance Security Guidance v4 and NIST SP 800-128, Guide for Security-Focused Configuration Management.

QUESTION NO: 29

What's the best way for organizations to establish a foundation for safeguarding data, upholding privacy, and meeting regulatory requirements in cloud applications?

A. By implementing end-to-end encryption and multi-factor authentication

B. By conducting regular security audits and updates

C. By deploying intrusion detection systems and monitoring

D. By integrating security at the architectural and design level

ANSWER: D

Explanation:

By integrating security at the architectural and design level is correct because security, privacy, and compliance are most effective when they are established as foundational requirements before a cloud application is built and deployed. This security-by-design approach incorporates data classification, identity and access-management design, least-privilege authorization, encryption requirements, secure interfaces, logging, resilience, and compliance controls into the application’s architecture. Establishing these decisions early enables controls to be applied consistently across components, environments, and releases, while also producing the design evidence and traceability needed to support regulatory obligations. It makes security a continuous engineering concern rather than a set of protections added after deployment. The Cloud Security Alliance emphasizes that cloud security architecture must account for governance, risk, compliance, data security, and application security in a coordinated way. Its Cloud Controls Matrix provides a control framework that organizations can use to translate those requirements into cloud architecture and assurance activities. Similarly, the CSA guidance on security guidance for critical areas of cloud computing treats application security and information governance as core architectural concerns. See the Cloud Controls Matrix and the CSA Security Guidance v4.

QUESTION NO: 30

What is the newer application development methodology and philosophy focused on automation of application development and deployment?

A. Agile

B. BusOps

C. DevOps

D. SecDevOps

E. Scrum

ANSWER: C

Explanation:

DevOps is the application-development methodology and organizational philosophy centered on bringing development and operations practices together to deliver software more rapidly, reliably, and repeatedly. A defining DevOps practice is automation across the software delivery lifecycle: source-code integration, builds, testing, infrastructure provisioning, releases, deployment, monitoring, and operational feedback. This automation is commonly implemented through continuous integration and continuous delivery/deployment pipelines, infrastructure as code, configuration management, and automated observability controls. Rather than treating deployment as a separate, manual handoff after development is complete, DevOps promotes shared responsibility and fast feedback between the people building applications and the people operating them. This directly matches a methodology focused on automating application development and deployment. In cloud environments, DevOps automation is especially important because infrastructure can be programmatically created, changed, validated, and deployed at scale. Cloud Security Alliance guidance also recognizes DevSecOps as an evolution that integrates security practices into these automated DevOps processes, but the foundational development-and-operations methodology described here is DevOps. See the AWS overview of DevOps and NIST's Strategies for the Integration of Software Supply Chain Security in DevSecOps CI/CD Pipelines.

QUESTION NO: 31

Which of the following best explains how Multifactor Authentication (MFA) helps prevent identity-based attacks?

A. MFA relies on physical tokens and biometrics to secure accounts.

B. MFA requires multiple forms of validation that an attacker would have to compromise.

C. MFA requires and uses more complex passwords to secure accounts.

D. MFA eliminates the need for passwords through single sign-on.

ANSWER: B

Explanation:

MFA requires multiple forms of validation that an attacker would have to compromise. This is the best explanation because MFA strengthens an authentication decision by requiring evidence from separate factor categories, such as something a user knows, something they possess, or something they are. A stolen password alone therefore is not normally sufficient to authenticate as the targeted user; the attacker must also satisfy the additional required factor or factors. This substantially raises the effort and reduces the likelihood of successful identity-based attacks that depend on compromised credentials, including password spraying, credential stuffing, and reuse of passwords exposed in a breach. The protection is strongest when the factors are genuinely independent and when organizations use phishing-resistant methods, such as FIDO2 security keys or platform authenticators, for higher-risk access. MFA should also be implemented as part of an identity and access management program that includes strong enrollment, recovery, logging, and monitoring controls. CSA's Cloud Controls Matrix identifies multifactor authentication as an important identity and access-management control, while CSA Security Guidance discusses IAM practices for protecting cloud identities. See the Cloud Controls Matrix and CSA Security Guidance v5.

QUESTION NO: 32

Without virtualization, there is no cloud.

A. False

B. True

ANSWER: A

Explanation:

False is correct because virtualization is a widely used enabling technology for cloud computing, but it is not a required defining characteristic of cloud services. The NIST definition identifies cloud computing through five essential characteristics: on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service. This definition is deliberately technology-neutral; it does not make virtual machines or hypervisors mandatory. A provider can deliver cloud characteristics using automated provisioning and orchestration of physical, bare-metal resources, containers, or other abstraction mechanisms. For example, bare-metal cloud offerings can provide self-service, API-based provisioning, metering, and elastic allocation while assigning customers dedicated physical servers rather than virtual machines. Virtualization remains important from a security and operations perspective because it can improve utilization, isolation, portability, and scalability, while also introducing risks involving the hypervisor and virtual network layers. However, those benefits do not change the fundamental distinction: cloud computing is defined by its service model and operational characteristics, not by a specific implementation technology. See the NIST Definition of Cloud Computing (SP 800-145) and the Cloud Security Alliance Security Guidance v4.

QUESTION NO: 33

What does Zero Trust Network Access (ZTNA) primarily use to control access to applications?

A. Geolocation data exclusively

B. Username and password

C. IP address and port number

D. Identity, device, and contextual factors

ANSWER: D

Explanation:

Identity, device, and contextual factors is correct because Zero Trust Network Access makes application-access decisions using multiple continuously evaluated signals rather than relying on a traditional network perimeter. A ZTNA service first establishes a strong identity signal, commonly through an identity provider and authentication controls such as multifactor authentication. It then evaluates the security posture of the requesting device, such as whether the device is managed, compliant, encrypted, and running required endpoint protections. Context supplies further risk information, including the application being requested, the user's role, location, time, network conditions, and observed session risk. Based on these signals and least-privilege policy, ZTNA grants access only to the specific authorized application or resource, rather than placing the user broadly on the internal network. This identity-centric, context-aware approach supports continuous verification and reduces opportunities for unauthorized access and lateral movement. NIST's Zero Trust Architecture describes access decisions as dynamically evaluated using identity, device state, and environmental or behavioral attributes. The Cloud Security Alliance likewise identifies zero trust as an approach centered on explicit verification and least-privilege access. See NIST SP 800-207, Zero Trust Architecture and the Cloud Security Alliance Zero Trust Advancement Center.

QUESTION NO: 34

Which activity is a critical part of the Post-Incident Analysis phase in cybersecurity incident response?

A. Notifying affected parties

B. Isolating affected systems

C. Restoring services to normal operations

D. Documenting lessons learned and improving future responses

ANSWER: D

Explanation:

Documenting lessons learned and improving future responses is a critical post-incident activity because the organization must turn the incident's technical, operational, and business findings into measurable improvements. After containment,

eradication, and recovery have taken place, the incident response team should conduct a structured review of what occurred, how the incident was detected and handled, which controls or processes worked, and where response delays or gaps appeared. This review should produce appropriately protected records, including the timeline, impact, root cause or contributing factors, decisions made, evidence of response effectiveness, and recommended corrective actions. Lessons learned should feed updates to incident-response playbooks, escalation paths, logging and monitoring coverage, training, security controls, and testing exercises. Assigning owners and target dates for these improvements helps ensure that findings result in lasting risk reduction rather than a report that is simply archived. CSA guidance identifies post-incident activities, including lessons learned and process improvement, as necessary to mature cloud incident-response capabilities. NIST likewise describes the lessons-learned process as an opportunity to improve policies, procedures, and technical capabilities after an incident. See the Cloud Security Alliance Cloud Controls Matrix and NIST SP 800-61 Rev. 2.

QUESTION NO: 35

What is a primary objective during the Detection and Analysis phase of incident response?

A. Developing and updating incident response policies

B. Validating alerts and estimating the scope of incidents

C. Performing detailed forensic investigations

D. Implementing network segmentation and isolation

ANSWER: B

Explanation:

Validating alerts and estimating the scope of incidents is the primary objective because Detection and Analysis turns potentially meaningful security signals into a confirmed, actionable understanding of an event. Security teams assess alerts from monitoring tools, logs, endpoint telemetry, cloud audit records, threat intelligence, and user reports to determine whether an actual incident has occurred. Once an incident is identified, responders collect and correlate relevant evidence to characterize its nature, affected cloud services or assets, suspected entry point, time frame, accounts involved, data exposure risk, and likely business impact. This initial scoping supports appropriate prioritization and gives incident leadership the information needed to make timely containment decisions. In a cloud environment, this commonly includes reviewing provider audit trails, identity activity, API calls, workload logs, and configuration changes. NIST's incident-handling guidance describes Detection and Analysis as the stage in which organizations identify and investigate precursors and indicators, then document and prioritize incidents according to their scope and impact. This evidence-based assessment ensures that subsequent response actions are proportionate and directed at the actual incident conditions. See NIST SP 800-61 Revision 2, Computer Security Incident Handling Guide.