



"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks.com

Warning: Keep connected with our support team for latest updates

Question: 1

What native runtime is Open Container Initiative (OCI) compliant?

- A. RunC
- B. RunV
- C. kata-containers
- D. gvisor

Answer: A

Explanation:

The Open Container Initiative (OCI) publishes open specifications for container images and container runtimes so that tools across the ecosystem remain interoperable. When a runtime is "OCI-compliant," it means it implements the OCI Runtime Specification (how to run a container from a filesystem bundle and configuration) and/or works cleanly with OCI image formats through the usual layers (image → unpack → runtime). runC is the best-known, widely used reference implementation of the OCI runtime specification and is the low-level runtime underneath many higher-level systems. In Kubernetes, you typically interact with a higher-level container runtime (such as containerd or CRI-O) through the Container Runtime Interface (CRI). That higher-level runtime then uses a low-level OCI runtime to actually create Linux namespaces/cgroups, set up the container process, and start it. In many default installations, containerd delegates to runC for this low-level "create/start" work.

The other options are related but differ in what they are: Kata Containers uses lightweight VMs to provide stronger isolation while still presenting a container-like workflow; gVisor provides a userspace kernel for sandboxing containers; these can be used with Kubernetes via compatible integrations, but the canonical "native OCI runtime" answer in most curricula is runC. Finally, "runV" is not a common modern Kubernetes runtime choice in typical OCI discussions. So the most correct, standards-based answer here is A (runC) because it directly implements the OCI runtime spec and is commonly used as the default low-level runtime behind CRI implementations.

Question: 2

Which API object is the recommended way to run a scalable, stateless application on your cluster?

- A. ReplicaSet
- B. Deployment
- C. DaemonSet
- D. Pod

Answer: B

Explanation:

For a scalable, stateless application, Kubernetes recommends using a Deployment because it provides a higher-level, declarative management layer over Pods. A Deployment doesn't just "run replicas"; it manages the entire lifecycle of rolling out new versions, scaling up/down, and recovering from failures by continuously reconciling the current cluster state to the desired state you define. Under the hood, a Deployment typically creates and manages a ReplicaSet, and that ReplicaSet ensures a specified number of Pod replicas are running at all times. This layering is the key: you get ReplicaSet's self-healing replica maintenance plus Deployment's rollout/rollback strategies and revision history.

Why not the other options? A Pod is the smallest deployable unit, but it's not a scalable controller— if a Pod dies, nothing automatically replaces it unless a controller owns it. A ReplicaSet can maintain N replicas, but it does not provide the full rollout orchestration (rolling updates, pause/resume, rollbacks, and revision tracking) that you typically want for stateless apps that ship frequent releases. A DaemonSet is for node-scoped workloads (one Pod per node or subset of nodes), like log shippers or node agents, not for "scale by replicas."

For stateless applications, the Deployment model is especially appropriate because individual replicas are interchangeable; the application does not require stable network identities or persistent storage per replica. Kubernetes can freely replace or reschedule Pods to maintain availability. Deployment strategies (like RollingUpdate) allow you to upgrade without downtime by gradually replacing old replicas with new ones while keeping the Service endpoints healthy. That combination—declarative desired state, self-healing, and controlled updates—makes Deployment the recommended object for scalable stateless workloads.

Question: 3

A CronJob is scheduled to run by a user every one hour. What happens in the cluster when it's time for this CronJob to run?

- A. Kubelet watches API Server for CronJob objects. When it's time for a Job to run, it runs the Pod directly.
- B. Kube-scheduler watches API Server for CronJob objects, and this is why it's called kube-scheduler.
- C. CronJob controller component creates a Pod and waits until it finishes to run.
- D. CronJob controller component creates a Job. Then the Job controller creates a Pod and waits until it finishes to run.

Answer: D

Explanation:

CronJobs are implemented through Kubernetes controllers that reconcile desired state. When the scheduled time arrives, the CronJob controller (part of the controller-manager set of control plane controllers) evaluates the CronJob object's schedule and determines whether a run should be started. Importantly, CronJob does not create Pods directly as its primary mechanism. Instead, it creates a Job object for each scheduled execution. That Job object then becomes the responsibility of the Job controller, which creates one or more Pods to complete the Job's work and monitors them until completion. This separation of concerns is why option D is correct.

This design has practical benefits. Jobs encapsulate "run-to-completion" semantics: retries, backoff limits, completion counts, and tracking whether the work has succeeded. CronJob focuses on the temporal triggering aspect (schedule, concurrency policy, starting deadlines, history limits), while Job focuses on the execution aspect (create Pods, ensure completion, retry on failure).

Option A is incorrect because kubelet is a node agent; it does not watch CronJob objects and doesn't decide when a schedule triggers. Kubelet reacts to Pods assigned to its node and ensures containers run there.

Option B is incorrect because kube-scheduler schedules Pods to nodes after they exist (or are created by controllers); it does not trigger CronJobs. Option C is incorrect because CronJob does not usually create a Pod and wait directly; it delegates via a Job, which then manages Pods and completion.

So, at runtime: CronJob controller creates a Job; Job controller creates the Pod(s); scheduler assigns those Pods to nodes; kubelet runs them; Job controller observes success/failure and updates status; CronJob controller manages run history and concurrency rules.

Question: 4

What is the purpose of the kubelet component within a Kubernetes cluster?

- A. A dashboard for Kubernetes clusters that allows management and troubleshooting of applications.
- B. A network proxy that runs on each node in your cluster, implementing part of the Kubernetes Service concept.
- C. A component that watches for newly created Pods with no assigned node, and selects a node for them to run on.
- D. An agent that runs on each node in the cluster. It makes sure that containers are running in a Pod.

Answer: D

Explanation:

The kubelet is the primary node agent in Kubernetes. It runs on every worker node (and often on control-plane nodes too if they run workloads) and is responsible for ensuring that containers described by PodSpecs are actually running and healthy on that node. The kubelet continuously watches the Kubernetes API (via the control plane) for Pods that have been scheduled to its node, then it collaborates with the node's container runtime (through CRI) to pull images, create containers, start them, and manage their lifecycle. It also mounts volumes, configures the Pod's networking (working with the CNI plugin), and reports Pod and node status back to the API server.

Option D captures the core: "an agent on each node that makes sure containers are running in a Pod." That includes executing probes (liveness, readiness, startup), restarting containers based on the Pod's restartPolicy, and enforcing resource constraints in coordination with the runtime and OS.

Why the other options are wrong: A describes the Kubernetes Dashboard (or similar UI tools), not kubelet. B describes kube-proxy, which programs node-level networking rules (iptables/ipvs/eBPF depending on implementation) to implement Service virtual IP behavior. C describes the kube-scheduler, which selects a node for Pods that do not yet have an assigned node.

A useful way to remember kubelet's role is: scheduler decides where, kubelet makes it happen there. Once the scheduler binds a Pod to a node, kubelet becomes responsible for reconciling "desired state" (PodSpec) with "observed state" (running containers). If a container crashes, kubelet will restart it according to policy; if an image is missing, it will pull it; if a Pod is deleted, it will stop

containers and clean up. This node-local reconciliation loop is fundamental to Kubernetes' selfhealing and declarative operation model.

Question: 5

What is the default value for authorization-mode in Kubernetes API server?

- A. --authorization-mode=RBAC
- B. --authorization-mode=AlwaysAllow
- C. --authorization-mode=AlwaysDeny
- D. --authorization-mode=ABAC

Answer: B

Explanation:

The Kubernetes API server supports multiple authorization modes that determine whether an authenticated request is allowed to perform an action (verb) on a resource. Historically, the API server's default authorization mode was AlwaysAllow, meaning that once a request was authenticated, it would be authorized without further checks. That is why the correct answer here is B.

However, it's crucial to distinguish "default flag value" from "recommended configuration." In production clusters, running with AlwaysAllow is insecure because it effectively removes authorization controls—any authenticated user (or component credential) could do anything the API permits. Modern Kubernetes best practices strongly recommend enabling RBAC (Role-Based Access Control), often alongside Node and Webhook authorization, so that permissions are granted explicitly using Roles/ClusterRoles and RoleBindings/ClusterRoleBindings. Many managed Kubernetes distributions and kubeadm-based setups commonly enable RBAC by default as part of cluster bootstrap profiles, even if the API server's historical default flag value is AlwaysAllow.

So, the exam-style interpretation of this question is about the API server flag default, not what most real clusters should run. With RBAC enabled, authorization becomes granular: you can control who can read Secrets, who can create Deployments, who can exec into Pods, and so on, scoped to namespaces or

cluster-wide. ABAC (Attribute-Based Access Control) exists but is generally discouraged compared to RBAC because it relies on policy files and is less ergonomic and less commonly used. AlwaysDeny is useful for hard lockdown testing but not for normal clusters.

In short: AlwaysAllow is the API server's default mode (answer B), but RBAC is the secure, recommended choice you should expect to see enabled in almost any serious Kubernetes environment.

Question: 6

Let's assume that an organization needs to process large amounts of data in bursts, on a cloud-based Kubernetes cluster. For instance: each Monday morning, they need to run a batch of 1000 compute jobs of 1 hour each, and these jobs must be completed by Monday night. What's going to be the most cost-effective method?

- A. Run a group of nodes with the exact required size to complete the batch on time, and use a combination of taints, tolerations, and nodeSelectors to reserve these nodes to the batch jobs.
- B. Leverage the Kubernetes Cluster Autoscaler to automatically start and stop nodes as they're needed.
- C. Commit to a specific level of spending to get discounted prices (with e.g. "reserved instances" or similar mechanisms).
- D. Use PriorityClasses so that the weekly batch job gets priority over other workloads running on the cluster, and can be completed on time.

Answer: B

Explanation:

Burst workloads are a classic elasticity problem: you need large capacity for a short window, then very little capacity the rest of the week. The most cost-effective approach in a cloud-based Kubernetes environment is to scale infrastructure dynamically, matching node count to current demand. That's exactly what Cluster Autoscaler is designed for: it adds nodes when Pods cannot be scheduled due to insufficient resources and removes nodes when they become underutilized and can be drained safely. Therefore B is correct.

Option A can work operationally, but it commonly results in paying for a reserved “standing army” of nodes that sit idle most of the week—wasteful for bursty patterns unless the nodes are repurposed for other workloads. Taints/tolerations and nodeSelectors are placement tools; they don’t reduce cost by themselves and may increase waste if they isolate nodes. Option D (PriorityClasses) affects which Pods get scheduled first given available capacity, but it does not create capacity. If the cluster doesn’t have enough nodes, high priority Pods will still remain Pending. Option C (reserved instances or committed-use discounts) can reduce unit price, but it assumes relatively predictable baseline usage. For true bursts, you usually want a smaller baseline plus autoscaling, and optionally combine it with discounted capacity types if your cloud supports them.

In Kubernetes terms, the control loop is: batch Jobs create Pods → scheduler tries to place Pods → if many Pods are Pending due to insufficient CPU/memory, Cluster Autoscaler observes this and increases the node group size → new nodes join and kube-scheduler places Pods → after jobs finish and nodes become empty, Cluster Autoscaler drains and removes nodes. This matches cloud-native principles: elasticity, pay-for-what-you-use, and automation. It minimizes idle capacity while still meeting the completion deadline.

Question: 7

What is a Kubernetes service with no cluster IP address called?

- A. Headless Service
- B. Nodeless Service
- C. IPLess Service
- D. Specless Service

Answer: A

Explanation:

A Kubernetes Service normally provides a stable virtual IP (ClusterIP) and a DNS name that loadbalances traffic across matching Pods. A headless Service is a special type of Service where Kubernetes does not allocate a ClusterIP. Instead, the Service’s DNS returns individual Pod IPs (or other endpoint records), allowing clients to connect directly to specific backends rather than through a single virtual IP. That is why the correct answer is A (Headless Service).

Headless Services are created by setting `spec.clusterIP: None`. When you do this, kube-proxy does not program load-balancing rules for a virtual IP because there isn't one. Instead, service discovery is handled via DNS records that point to the actual endpoints. This behavior is especially important for stateful or identity-sensitive systems where clients must talk to a particular replica (for example, databases, leader/follower clusters, or StatefulSet members).

This is also why headless Services pair naturally with StatefulSets. StatefulSets provide stable network identities (pod-0, pod-1, etc.) and stable DNS names. The headless Service provides the DNS domain that resolves each Pod's stable hostname to its IP, enabling peer discovery and consistent addressing even as Pods move between nodes.

The other options are distractors: "Nodeless," "IPLess," and "Specless" are not Kubernetes Service types. In the core API, the Service "types" are things like ClusterIP, NodePort, LoadBalancer, and ExternalName; "headless" is a behavioral mode achieved through the ClusterIP field.

In short: a headless Service removes the virtual IP abstraction and exposes endpoint-level discovery. It's a deliberate design choice when load-balancing is not desired or when the application itself handles routing, membership, or sharding.

Question: 8

CI/CD stands for:

- A. Continuous Information / Continuous Development
- B. Continuous Integration / Continuous Development
- C. Cloud Integration / Cloud Development
- D. Continuous Integration / Continuous Deployment

Answer: D

Explanation:

CI/CD is a foundational practice for delivering software rapidly and reliably, and it maps strongly to cloud native delivery workflows commonly used with Kubernetes. CI stands for Continuous Integration: developers merge code changes frequently into a shared repository, and automated systems build and test those changes

to detect issues early. CD is commonly used to mean Continuous Delivery or Continuous Deployment depending on how far automation goes. In many certification contexts and simplified definitions like this question, CD is interpreted as Continuous Deployment, meaning every change that passes the automated pipeline is automatically released to production. That matches option D.

In a Kubernetes context, CI typically produces artifacts such as container images (built from Dockerfiles or similar build definitions), runs unit/integration tests, scans dependencies, and pushes images to a registry. CD then promotes those images into environments by updating Kubernetes manifests (Deployments, Helm charts, Kustomize overlays, etc.). Progressive delivery patterns (rolling updates, canary, blue/green) often use Kubernetes-native controllers and Service routing to reduce risk.

Why the other options are incorrect: "Continuous Development" isn't the standard "D" term; it's ambiguous and not the established acronym expansion. "Cloud Integration/Cloud Development" is unrelated. Continuous Delivery (in the stricter sense) means changes are always in a deployable state and releases may still require a manual approval step, while Continuous Deployment removes that final manual gate. But because the option set explicitly includes "Continuous Deployment," and that is one of the accepted canonical expansions for CD, D is the correct selection here.

Practically, CI/CD complements Kubernetes' declarative model: pipelines update desired state (Git or manifests), and Kubernetes reconciles it. This combination enables frequent releases, repeatability, reduced human error, and faster recovery through automated rollbacks and controlled rollout strategies.

Question: 9

What default level of protection is applied to the data in Secrets in the Kubernetes API?

- A. The values use AES symmetric encryption
- B. The values are stored in plain text
- C. The values are encoded with SHA256 hashes
- D. The values are base64 encoded

Answer: D

Explanation:

Kubernetes Secrets are designed to store sensitive data such as tokens, passwords, or certificates and make

them available to Pods in controlled ways (as environment variables or mounted files).

However, the default protection applied to Secret values in the Kubernetes API is base64 encoding, not encryption. That is why D is correct. Base64 is an encoding scheme that converts binary data into ASCII text; it is reversible and does not provide confidentiality.

By default, Secret objects are stored in the cluster's backing datastore (commonly etcd) as base64- encoded strings inside the Secret manifest. Unless the cluster is configured for encryption at rest, those values are effectively stored unencrypted in etcd and may be visible to anyone who can read etcd directly or who has API permissions to read Secrets. This distinction is critical for security: base64 can prevent accidental issues with special characters in YAML/JSON, but it does not protect against attackers.

Option A is only correct if encryption at rest is explicitly configured on the API server using an EncryptionConfiguration (for example, AES-CBC or AES-GCM providers). Many managed Kubernetes offerings enable encryption at rest for etcd as an option or by default, but that is a deployment choice, not the universal Kubernetes default. Option C is incorrect because hashing is used for verification, not for secret retrieval; you typically need to recover the original value, so hashing isn't suitable for Secrets. Option B ("plain text") is misleading: the stored representation is base64- encoded, but because base64 is reversible, the security outcome is close to plain text unless encryption at rest and strict RBAC are in place.

The correct operational stance is: treat Kubernetes Secrets as sensitive; lock down access with RBAC, enable encryption at rest, avoid broad Secret read permissions, and consider external secret managers when appropriate. But strictly for the question's wording—default level of protection— base64 encoding is the right answer.

Question: 10

What function does kube-proxy provide to a cluster?

- A. Implementing the Ingress resource type for application traffic.
- B. Forwarding data to the correct endpoints for Services.
- C. Managing data egress from the cluster nodes to the network.
- D. Managing access to the Kubernetes API.

Answer: B

Explanation:

kube-proxy is a node-level networking component that helps implement the Kubernetes Service abstraction. Services provide a stable virtual IP and DNS name that route traffic to a set of Pods (endpoints). kube-proxy watches the API for Service and EndpointSlice/Endpoints changes and then programs the node's networking rules so that traffic sent to a Service is forwarded (load-balanced) to one of the correct backend Pod IPs. This is why B is correct.

Conceptually, kube-proxy turns the declarative Service configuration into concrete dataplane behavior. Depending on the mode, it may use iptables rules, IPVS, or integrate with eBPF-capable networking stacks (sometimes kube-proxy is replaced or bypassed by CNI implementations, but the classic kube-proxy role remains the canonical answer). In iptables mode, kube-proxy creates NAT rules that rewrite traffic from the Service virtual IP to one of the Pod endpoints. In IPVS mode, it programs kernel load-balancing tables for more scalable service routing. In all cases, the job is to connect "Service IP/port" to "Pod IP/port endpoints."

Option A is incorrect because Ingress is a separate API resource and requires an Ingress Controller (like NGINX Ingress, HAProxy, Traefik, etc.) to implement HTTP routing, TLS termination, and host/path rules. kube-proxy is not an Ingress controller. Option C is incorrect because general node egress management is not kube-proxy's responsibility; egress behavior typically depends on the CNI plugin, NAT configuration, and network policies. Option D is incorrect because API access control is handled by the API server's authentication/authorization layers (RBAC, webhooks, etc.), not kube-proxy.

So kube-proxy's essential function is: keep node networking rules in sync so that Service traffic reaches the right Pods. It is one of the key components that makes Services "just work" across nodes without clients needing to know individual Pod IPs.

Question: 11

How long should a stable API element in Kubernetes be supported (at minimum) after deprecation?

- A. 9 months
- B. 24 months
- C. 12 months
- D. 6 months

Answer: C

Explanation:

Kubernetes has a formal API deprecation policy to balance stability for users with the ability to evolve the platform. For a stable (GA) API element, Kubernetes commits to supporting that API for a minimum period after it is deprecated. The correct minimum in this question is 12 months, which corresponds to option C.

In practice, Kubernetes releases occur roughly every three to four months, and the deprecation policy is commonly communicated in terms of “releases” as well as time. A GA API that is deprecated in one release is typically kept available for multiple subsequent releases, giving cluster operators and application teams time to migrate manifests, client libraries, controllers, and automation. This matters because Kubernetes is often at the center of production delivery pipelines; abrupt API removals would break deployments, upgrades, and tooling. By guaranteeing a minimum support window, Kubernetes enables predictable upgrades and safer lifecycle management.

This policy also encourages teams to track API versions and plan migrations. For example, workloads might start on a beta API (which can change), but once an API reaches stable, users can expect a stronger compatibility promise. Deprecation warnings help surface risk early. In many clusters, you’ll see API server warnings and tooling hints when manifests use deprecated fields/versions, allowing

proactive remediation before the removal release.

Options 6 or 9 months would be too short for many enterprises to coordinate changes across multiple teams and environments. 24 months may be true for some ecosystems, but the Kubernetes stated minimum in this exam-style framing is 12 months. The key operational takeaway is: don't ignore deprecation notices—they're your clock for migration planning. Treat API version upgrades as part of routine cluster lifecycle hygiene to avoid being blocked during Kubernetes version upgrades when deprecated APIs are finally removed.

Question: 12

What is the name of the lightweight Kubernetes distribution built for IoT and edge computing?

- A. OpenShift
- B. k3s
- C. RKE
- D. k1s

Answer: B

Explanation:

Edge and IoT environments often have constraints that differ from traditional datacenters: limited CPU/RAM, intermittent connectivity, smaller footprints, and a desire for simpler operations. k3s is a well-known lightweight Kubernetes distribution designed specifically to run in these environments, making B the correct answer.

What makes k3s "lightweight" is that it packages Kubernetes components in a simplified way and reduces operational overhead. It typically uses a single binary distribution and can run with an embedded datastore option for smaller installations (while also supporting external datastores for HA use cases). It streamlines dependencies and is aimed at faster installation and reduced resource consumption, which is ideal for edge nodes, IoT gateways, small servers, labs, and development environments.

By contrast, OpenShift is a Kubernetes distribution focused on enterprise platform capabilities, with additional security defaults, integrated developer tooling, and a larger operational footprint—excellent for many enterprises but not "built for IoT and edge" as the defining characteristic. RKE (Rancher Kubernetes Engine) is a Kubernetes installer/engine used to deploy Kubernetes, but it's not specifically the lightweight edge-focused distribution in the way k3s is. "k1s" is not a

standard, widely recognized Kubernetes distribution name in this context.

From a cloud native architecture perspective, edge Kubernetes distributions extend the same declarative and API-driven model to places where you want consistent operations across cloud, datacenter, and edge. You can apply GitOps patterns, standard manifests, and Kubernetes-native controllers across heterogeneous footprints. k3s provides that familiar Kubernetes experience while optimizing for constrained environments, which is why it has become a common choice for edge/IoT Kubernetes deployments.

Question: 13

Kubernetes allows you to automatically manage the number of nodes in your cluster to meet demand.

- A. Node Autoscaler
- B. Cluster Autoscaler
- C. Horizontal Pod Autoscaler
- D. Vertical Pod Autoscaler

Answer: B

Explanation:

Kubernetes supports multiple autoscaling mechanisms, but they operate at different layers. The question asks specifically about automatically managing the number of nodes in the cluster, which is the role of the Cluster Autoscaler—therefore B is correct.

Cluster Autoscaler monitors the scheduling state of the cluster. When Pods are pending because there are not enough resources (CPU/memory) available on existing nodes—meaning the scheduler cannot place them—Cluster Autoscaler can request that the underlying infrastructure (typically a cloud provider node group / autoscaling group) add nodes. Conversely, when nodes are underutilized and Pods can be rescheduled elsewhere, Cluster Autoscaler can drain those nodes (respecting disruption constraints like PodDisruptionBudgets) and then remove them to reduce cost. This aligns with cloud-native elasticity: scale infrastructure up and down automatically based on workload needs.

The other options are different: Horizontal Pod Autoscaler (HPA) changes the number of Pod replicas for a workload (like a Deployment) based on metrics (CPU utilization, memory, or custom metrics). It scales the application layer, not the

node layer. Vertical Pod Autoscaler (VPA) changes resource requests/limits (CPU/memory) for Pods, effectively “scaling up/down” the size of individual Pods. It also does not directly change node count, though its adjustments can influence scheduling pressure. “Node Autoscaler” is not the canonical Kubernetes component name used in standard terminology; the widely referenced upstream component for node count is Cluster Autoscaler.

In real systems, these autoscalers often work together: HPA increases replicas when traffic rises; that may cause Pods to go Pending if nodes are full; Cluster Autoscaler then adds nodes; scheduling proceeds; later, traffic drops, HPA reduces replicas and Cluster Autoscaler removes nodes. This layered approach provides both performance and cost efficiency.

Question: 14

Which of the following statements is correct concerning Open Policy Agent (OPA)?

- A. The policies must be written in Python language.
- B. Kubernetes can use it to validate requests and apply policies.
- C. Policies can only be tested when published.
- D. It cannot be used outside Kubernetes.

Answer: B

Explanation:

Open Policy Agent (OPA) is a general-purpose policy engine used to define and enforce policy across different systems. In Kubernetes, OPA is commonly integrated through admission control (often via Gatekeeper or custom admission webhooks) to validate and/or mutate requests before they are persisted in the cluster. This makes B correct: Kubernetes can use OPA to validate API requests and apply policy decisions.

Kubernetes’ admission chain is where policy enforcement naturally fits. When a user or controller submits a request (for example, to create a Pod), the API server can call external admission webhooks. Those webhooks can evaluate the request against policy—such as “no privileged containers,” “images must come from approved registries,” “labels must include cost-center,” or “Ingress must enforce TLS.” OPA’s policy language (Rego) allows expressing these rules in a declarative form, and the decision (“allow/deny” and sometimes patches) is returned to the API server. This enforces governance consistently and centrally.

Option A is incorrect because OPA policies are written in Rego, not Python. Option C is incorrect because policies can be tested locally and in CI pipelines before deployment; in fact, testability is a key advantage. Option D is incorrect because OPA is designed to be platform-agnostic—it can be used with APIs, microservices, CI/CD pipelines, service meshes, and infrastructure tools, not only Kubernetes.

From a Kubernetes fundamentals view, OPA complements RBAC: RBAC answers “who can do what to which resources,” while OPA-style admission policies answer “even if you can create this resource, does it meet our organizational rules?” Together they help implement defense in depth: authentication + authorization + policy admission + runtime security controls. That is why OPA is widely used to enforce security and compliance requirements in Kubernetes environments.

Question: 15

In a cloud native world, what does the IaC abbreviation stand for?

- A. Infrastructure and Code
- B. Infrastructure as Code
- C. Infrastructure above Code
- D. Infrastructure across Code

Answer: B

Explanation:

IaC stands for Infrastructure as Code, which is option B. In cloud native environments, IaC is a core operational practice: infrastructure (networks, clusters, load balancers, IAM roles, storage classes, DNS records, and more) is defined using code-like, declarative configuration rather than manual, click-driven changes. This approach mirrors Kubernetes' own declarative model—where you define desired state in manifests and controllers reconcile the cluster to match.

IaC improves reliability and velocity because it makes infrastructure repeatable, version-controlled, reviewable, and testable. Teams can store infrastructure definitions in Git, use pull requests for change review, and run automated checks to validate formatting, policies, and safety constraints. If an environment must be recreated (disaster recovery, test environments, regional expansion), IaC enables consistent reproduction with fewer human errors.

In Kubernetes-centric workflows, IaC often covers both the base platform and the workloads layered on top. For example, provisioning might include the Kubernetes control plane, node pools, networking, and identity integration,

while Kubernetes manifests (or Helm/Kustomize) define Deployments, Services, RBAC, Ingress, and storage resources.

GitOps extends this further by continuously reconciling cluster configuration from a Git source of truth.

The incorrect options (Infrastructure and Code / above / across) are not standard terms. The key idea is “infrastructure treated like software”: changes are made through code commits, go through CI checks, and are rolled out in controlled ways. This aligns with cloud native goals: faster iteration, safer operations, and easier auditing. In short, IaC is the operational backbone that makes Kubernetes and cloud platforms manageable at scale, enabling consistent environments and reducing configuration drift.

You’re right — my previous 16–30 were not taken from your PDF. Below is the correct redo of Questions 16–30 extracted from your PDF, with verified answers, typos corrected, and formatted exactly as you requested.

Question: 16

In which framework do the developers no longer have to deal with capacity, deployments, scaling and fault tolerance, and OS?

- A. Docker Swarm
- B. Kubernetes
- C. Mesos
- D. Serverless

Answer: D

Explanation:

Serverless is the model where developers most directly avoid managing server capacity, OS operations, and much of the deployment/scaling/fault-tolerance mechanics, which is why D is correct. In serverless computing (commonly Function-as-a-Service, FaaS, and managed serverless container platforms), the provider abstracts away the underlying servers. You typically deploy code (functions) or a container image, define triggers (HTTP events, queues, schedules), and the platform automatically provisions the required compute, scales it based on demand, and handles much of the availability and fault tolerance behind the scenes.

It’s important to compare this to Kubernetes: Kubernetes does automate scheduling, self-healing, rolling updates, and

scaling, but it still requires you (or your platform team) to design and operate cluster capacity, node pools, upgrades, runtime configuration, networking, and baseline reliability controls. Even in managed Kubernetes services, you still choose node sizes, scale policies, and operational configuration. Kubernetes reduces toil, but it does not eliminate infrastructure concerns in the same way serverless does.

Docker Swarm and Mesos are orchestration platforms that schedule workloads, but they also require managing the underlying capacity and OS-level aspects. They are not “no longer have to deal with capacity and OS” frameworks.

From a cloud native viewpoint, serverless is about consuming compute as an on-demand utility. Kubernetes can be a foundation for a serverless experience (for example, with event-driven autoscaling or serverless frameworks), but the pure framework that removes the most operational burden from developers is serverless.

Question: 17

Which of the following characteristics is associated with container orchestration?

- A. Application message distribution
- B. Dynamic scheduling
- C. Deploying application JAR files
- D. Virtual machine distribution

Answer: B

Explanation:

A core capability of container orchestration is dynamic scheduling, so B is correct. Orchestration platforms (like Kubernetes) are responsible for deciding where containers (packaged as Pods in Kubernetes) should run, based on real-time cluster conditions and declared requirements. "Dynamic" means the system makes placement decisions continuously as workloads are created, updated, or fail, and as cluster capacity changes.

In Kubernetes, the scheduler evaluates Pods that have no assigned node, filters nodes that don't meet requirements (resources, taints/tolerations, affinity/anti-affinity, topology constraints), and then scores remaining nodes to pick the best target. This scheduling happens at runtime and adapts to the current state of the cluster. If nodes go down or Pods crash, controllers create replacements and the scheduler places them again—another aspect of dynamic orchestration.

The other options don't define container orchestration: "application message distribution" is more about messaging systems or service communication patterns, not orchestration. "Deploying application JAR files" is a packaging/deployment detail relevant to Java apps but not a defining orchestration capability. "Virtual machine distribution" refers to VM management rather than container orchestration; Kubernetes focuses on containers and Pods (even if those containers sometimes run in lightweight VMs via sandbox runtimes).

So, the defining trait here is that an orchestrator automatically and continuously schedules and reschedules workloads, rather than relying on static placement decisions.

Question: 18

Which of the following workload requires a headless Service while deploying into the namespace?

- A. StatefulSet
- B. CronJob
- C. Deployment
- D. DaemonSet

Answer: A

Explanation:

A StatefulSet commonly requires a headless Service, so A is the correct answer. In Kubernetes, StatefulSets are designed for workloads that need stable identities, stable network names, and often stable storage per replica. To support that stable identity model, Kubernetes typically uses a headless Service (spec.clusterIP: None) to provide DNS records that map directly to each Pod, rather than load-balancing behind a single virtual ClusterIP.

With a headless Service, DNS queries return individual endpoint records (the Pod IPs) so that each StatefulSet Pod can be addressed predictably, such as pod-0.service-name.namespace.svc.cluster.local. This is critical for clustered databases, quorum systems, and leader/follower setups where members must discover and address specific peers. The StatefulSet controller then ensures ordered creation/deletion and preserves identity (pod-0, pod-1, etc.), while the headless Service provides discovery for those stable hostnames.

CronJobs run periodic Jobs and don't require stable DNS identity for multiple replicas. Deployments manage stateless replicas and normally use a standard Service that load-balances across Pods. DaemonSets run one Pod per node, and while they can be exposed by Services, they do not intrinsically require headless discovery.

So while you can use a headless Service for other designs, StatefulSet is the workload type most associated with "requires a headless Service" due to how stable identities and per-Pod addressing work in Kubernetes.

Question: 18

What is Helm?

- A. An open source dashboard for Kubernetes.
- B. A package manager for Kubernetes applications.
- C. A custom scheduler for Kubernetes.

D. An end-to-end testing project for Kubernetes applications.

Answer: B

Explanation:

Helm is best described as a package manager for Kubernetes applications, making B correct. Helm packages Kubernetes resource manifests (Deployments, Services, ConfigMaps, Ingress, RBAC, etc.) into a unit called a chart. A chart includes templates and default values, allowing teams to parameterize deployments for different environments (dev/stage/prod) without rewriting YAML.

From an application delivery perspective, Helm solves common problems: repeatable installation, upgrade management, versioning, and sharing of standardized application definitions. Instead of copying and editing raw YAML, users install a chart and supply a values.yaml file (or CLI overrides) to configure image tags, replica counts, ingress hosts, resource requests, and other settings. Helm then renders templates into concrete Kubernetes manifests and applies them to the cluster.

Helm also manages releases: it tracks what has been installed and supports upgrades and rollbacks. This aligns with cloud native delivery practices where deployments are automated, reproducible, and auditable. Helm is commonly integrated into CI/CD pipelines and GitOps workflows (sometimes with charts stored in Git or Helm repositories).

The other options are incorrect: a dashboard is a UI like Kubernetes Dashboard; a scheduler is kube-scheduler (or custom scheduler implementations, but Helm is not that); end-to-end testing projects exist in the ecosystem, but Helm's role is packaging and lifecycle management of Kubernetes app definitions.

So the verified, standard definition is: Helm = Kubernetes package manager.

Question: 18

Which is the correct kubectl command to display logs in real time?

- A. `kubectl logs -p test-container-1`
- B. `kubectl logs -c test-container-1`
- C. `kubectl logs -l test-container-1`
- D. `kubectl logs -f test-container-1`

Answer: D

Explanation:

To stream logs in real time with kubectl, you use the follow option -f, so D is correct. In Kubernetes, kubectl logs retrieves logs from containers in a Pod. By default, it returns the current log output and exits. When you add -f, kubectl keeps the connection open and continuously prints new log lines as they are produced, similar to tail -f on Linux. This is especially useful for debugging live behavior, watching startup sequences, or monitoring an application during a rollout.

The other flags serve different purposes. -p (as seen in option A) requests logs from the previous instance of a container (useful after a restart/crash), not real-time streaming. -c (option B) selects a specific container within a multi-container Pod; it doesn't stream by itself (though it can be combined with -f). -l (option C) is used with kubectl logs to select Pods by label, but again it is not the streaming flag; streaming requires -f.

In real troubleshooting, you commonly combine flags, e.g. kubectl logs -f pod-name -c containername for streaming logs from a specific container, or kubectl logs -f -l app=myapp to stream from Pods matching a label selector (depending on kubectl behavior/version). But the key answer to "display logs in real time" is the follow flag: -f.

Therefore, the correct selection is D.

Question: 21

How to load and generate data required before the Pod startup?

- A. Use an init container with shared file storage.
- B. Use a PVC volume.
- C. Use a sidecar container with shared volume.
- D. Use another Pod with a PVC.

Answer: A

Explanation:

The Kubernetes-native mechanism to run setup steps before the main application containers start is an init container, so A is correct. Init containers run sequentially and must complete successfully before the regular containers in the Pod are started. This makes them ideal for preparing configuration, downloading artifacts, performing migrations, generating files, or waiting for dependencies.

The question specifically asks how to “load and generate data required before Pod startup.” The most common pattern is: an init container writes files into a shared volume (like an emptyDir volume) mounted by both the init container and the app container. When the init container finishes, the app container starts and reads the generated files. This is deterministic and aligns with Kubernetes Pod lifecycle semantics.

A sidecar container (option C) runs concurrently with the main container, so it is not guaranteed to complete work before startup. Sidecars are great for ongoing concerns (log shipping, proxies, config reloaders), but they are not the primary “before startup” mechanism. A PVC volume (option B) is just storage; it doesn’t itself perform generation or ensure ordering. “Another Pod with a PVC” (option D) introduces coordination complexity and still does not guarantee the data is prepared before this Pod starts unless you build additional synchronization.

Init containers are explicitly designed for this kind of pre-flight work, and Kubernetes guarantees ordering: all init containers complete in order, then the app containers begin. That guarantee is why A is the best and verified answer.

Question: 22

What is the core functionality of GitOps tools like Argo CD and Flux?

- A. They track production changes made by a human in a Git repository and generate a human-readable audit trail.
- B. They replace human operations with an agent that tracks Git commands.
- C. They automatically create pull requests when dependencies are outdated.
- D. They continuously compare the desired state in Git with the actual production state and notify or act upon differences.

Answer: D

Explanation:

The defining capability of GitOps controllers such as Argo CD and Flux is continuous reconciliation: they compare the desired state stored in Git to the actual state in the cluster and then alert and/or correct drift, making D correct. In

GitOps, Git becomes the single source of truth for declarative configuration (Kubernetes manifests, Helm charts, Kustomize overlays). The controller watches Git for changes and applies them, and it also watches the cluster for divergence.

This is more than “auditing human changes” (option A). GitOps does provide auditability because changes are made via commits and pull requests, but the core functionality is the reconciliation loop that keeps cluster state aligned with Git, including optional automated sync/remediation. Option B is not accurate because GitOps is not about tracking user Git commands; it’s about reconciling desired state definitions. Option C (automatically creating pull requests for outdated dependencies) is a useful feature in some tooling ecosystems, but it is not the central defining behavior of GitOps controllers.

In Kubernetes delivery terms, this approach improves reliability: rollouts become repeatable, configuration drift is detected, and recovery is simpler (reapply known-good state from Git). It also supports separation of duties: platform teams can control policies and base layers, while app teams propose changes via PRs.

So the verified statement is: GitOps tools continuously reconcile Git desired state with cluster actual state—exactly option D.

Question: 23

Which Kubernetes resource workload ensures that all (or some) nodes run a copy of a Pod?

- A. DaemonSet
- B. StatefulSet
- C. kubectl
- D. Deployment

Answer: A

Explanation:

A DaemonSet is the workload controller that ensures a Pod runs on all nodes or on a selected subset of nodes, so A is correct. DaemonSets are used for node-level agents and infrastructure components that must be present everywhere—examples include log collectors, monitoring agents, storage daemons, CNI components, and node security tools.

The DaemonSet controller watches for node additions/removals. When a new node joins the cluster, Kubernetes

automatically schedules a new DaemonSet Pod onto that node (subject to constraints such as node selectors, affinities, and taints/tolerations). When a node is removed, its DaemonSet Pod naturally disappears with it. This creates the “one per node” behavior that differentiates DaemonSets from other workload types.

A Deployment manages a replica count across the cluster, not “one per node.” A StatefulSet manages stable identity and ordered operations for stateful replicas; it does not inherently map one Pod to every node. kubectl is a CLI tool and not a workload resource.

DaemonSets can also be scoped: by using node selectors, node affinity, and tolerations, you can ensure Pods run only on GPU nodes, only on Linux nodes, only in certain zones, or only on nodes with a particular label. That’s why the question says “all (or some) nodes.”

Therefore, the correct and verified answer is DaemonSet (A).

Question: 24

What is CRD?

- A. Custom Resource Definition
- B. Custom Restricted Definition
- C. Customized RUST Definition
- D. Custom RUST Definition

Answer: A

Explanation:

A CRD is a CustomResourceDefinition, making A correct. Kubernetes is built around an API-driven model: resources like Pods, Services, and Deployments are all objects served by the Kubernetes API. CRDs allow you to extend the Kubernetes API by defining your own resource types. Once a CRD is installed, the API server can store and serve custom objects (Custom Resources) of that new type, and Kubernetes tooling (kubectl, RBAC, admission, watch mechanisms) can interact with them just like built-in resources.

CRDs are a core building block of the Kubernetes ecosystem because they enable operators and platform extensions. A typical pattern is: define a CRD that represents the desired state of some higher-level concept (for example, a database cluster, a certificate request, an application release), and then run a controller (often called an “operator”) that watches those custom resources and reconciles the cluster to match. That controller may create Deployments, StatefulSets,

Services, Secrets, or cloud resources to implement the desired state encoded in the custom resource.

The incorrect answers are made-up expansions. CRDs are not related to Rust in Kubernetes terminology, and “custom restricted definition” is not the standard meaning.

So the verified meaning is: CRD = CustomResourceDefinition, used to extend Kubernetes APIs and enable Kubernetes-native automation via controllers/operators.

Question: 25

The Kubernetes project work is carried primarily by SIGs. What does SIG stand for?

- A. Special Interest Group
- B. Software Installation Guide
- C. Support and Information Group
- D. Strategy Implementation Group

Answer: A

Explanation:

In Kubernetes governance and project structure, SIG stands for Special Interest Group, so A is correct. Kubernetes is a large open source project under the Cloud Native Computing Foundation (CNCF), and its work is organized into groups that focus on specific domains—such as networking, storage, node, scheduling, security, docs, testing, and many more. SIGs provide a scalable way to coordinate contributors, prioritize work, review design proposals (KEPs), triage issues, and manage releases in their area.

Each SIG typically has regular meetings, mailing lists, chat channels, and maintainers who guide the direction of that part of the project. For example, SIG Network focuses on Kubernetes networking architecture and components, SIG Storage on storage APIs and CSI integration, and SIG Scheduling on scheduler behavior and extensibility. This structure helps Kubernetes evolve while maintaining quality, review rigor, and community-driven decision making.

The other options are not part of Kubernetes project terminology. “Software Installation Guide” and the others might sound plausible, but they are not how Kubernetes defines SIGs.

Understanding SIGs matters operationally because many Kubernetes features and design changes originate from SIGs.

When you read Kubernetes enhancement proposals, release notes, or documentation, you’ll often see SIG ownership

and references. In short, SIGs are the primary organizational units for Kubernetes engineering and stewardship, and SIG = Special Interest Group.

Question: 26

What is the order of 4C's in Cloud Native Security, starting with the layer that a user has the most control over?

- A. Cloud -> Container -> Cluster -> Code
- B. Container -> Cluster -> Code -> Cloud
- C. Cluster -> Container -> Code -> Cloud
- D. Code -> Container -> Cluster -> Cloud

Answer: D

Explanation:

The Cloud Native Security “4C’s” model is commonly presented as Code, Container, Cluster, Cloud, ordered from the layer you control most directly to the one you control least—therefore D is correct. The idea is defense-in-depth across layers, recognizing that responsibilities are shared between developers, platform teams, and cloud providers.

Code is where users have the most direct control: application logic, dependencies, secure coding practices, secrets handling patterns, and testing. This includes validating inputs, avoiding vulnerabilities, and scanning dependencies. Next is the Container layer: building secure images, minimizing image size/attack surface, using non-root users, setting file permissions, and scanning images for known CVEs. Container security is about ensuring the artifact you run is trustworthy and hardened.

Then comes the Cluster layer: Kubernetes configuration and runtime controls, including RBAC, admission policies (OPA/Gatekeeper), Pod Security standards, network policies, runtime security, audit logging, and node hardening practices. Cluster controls determine what can run and how workloads interact. Finally, the Cloud layer includes the infrastructure and provider controls—IAM, VPC/networking, KMS, managed control plane protections, and physical security—which users influence through configuration but do not fully own.

The model’s value is prioritization: start with what you control most (code), then harden the container artifact, then enforce cluster policy and runtime protections, and finally ensure cloud controls are configured properly. This layered

approach aligns well with Kubernetes security guidance and modern shared-responsibility models.

Question: 27

Which group of container runtimes provides additional sandboxed isolation and elevated security?

- A. runc, cgroups
- B. docker, containerd
- C. runc, kata
- D. crun, cri-o

Answer: C

Explanation:

The runtimes most associated with sandboxed isolation are gVisor's runc and Kata Containers, making C correct.

Standard container runtimes (like containerd with runc) rely primarily on Linux namespaces and cgroups for isolation.

That isolation is strong for many use cases, but it shares the host kernel, which can be a concern for multi-tenant or high-risk workloads.

gVisor (runc) provides a user-space kernel-like layer that intercepts and mediates system calls, reducing the container's direct interaction with the host kernel. Kata Containers takes a different approach: it runs containers inside lightweight virtual machines, providing hardware-virtualization boundaries (or VM-like isolation) while still integrating into container workflows. Both are used to increase isolation compared to traditional containers, and both can be integrated with Kubernetes through compatible CRI/runtime configurations.

The other options are incorrect for the question's intent. "runc, cgroups" is not a meaningful pairing here (cgroups is a Linux resource mechanism, not a runtime). "docker, containerd" are commonly used container platforms/runtimes but are not specifically the "sandboxed isolation" category (containerd typically uses runc for standard isolation). "crun, cri-o" represents a low-level OCI runtime (crun) and a CRI implementation (CRI-O), again not specifically a sandboxed-isolation grouping.

So, when the question asks for the group that provides additional sandboxing and elevated security, the correct, well-established answer is runc + Kata.

Question: 28

What is the common standard for Service Meshes?

- A. Service Mesh Specification (SMS)
- B. Service Mesh Technology (SMT)
- C. Service Mesh Interface (SMI)
- D. Service Mesh Function (SMF)

Answer: C

Explanation:

A widely referenced interoperability standard in the service mesh ecosystem is the Service Mesh Interface (SMI), so C is correct. SMI was created to provide a common set of APIs for basic service mesh capabilities—helping users avoid being locked into a single mesh implementation for core features. While service meshes differ in architecture and implementation (e.g., Istio, Linkerd, Consul), SMI aims to standardize how common behaviors are expressed.

In cloud native architecture, service meshes address cross-cutting concerns for service-to-service communication: traffic policies, observability, and security (mTLS, identity). Rather than baking these concerns into every application, a mesh typically introduces data-plane proxies and a control plane to manage policy and configuration. SMI sits above those implementations as a common API model.

The other options are not commonly used industry standards. You may see other efforts and emerging APIs, but among the listed choices, SMI is the recognized standard name that appears in cloud native discussions and tooling integrations.

Also note a practical nuance: even with SMI, not every mesh implements every SMI spec fully, and many users still adopt mesh-specific CRDs and APIs for advanced features. But for this question's framing—"common standard"—Service Mesh Interface is the correct answer.

Question: 29

Which statement about Ingress is correct?

- A. Ingress provides a simple way to track network endpoints within a cluster.
- B. Ingress is a Service type like NodePort and ClusterIP.
- C. Ingress is a construct that allows you to specify how a Pod is allowed to communicate.
- D. Ingress exposes routes from outside the cluster to Services in the cluster.

Answer: D

Explanation:

Ingress is the Kubernetes API resource for defining external HTTP/HTTPS routing into the cluster, so D is correct. An Ingress object specifies rules such as hostnames (e.g., `app.example.com`), URL paths (e.g., `/api`), and TLS configuration, mapping those routes to Kubernetes Services. This provides Layer 7 routing capabilities beyond what a basic Service offers.

Ingress is not a Service type (so B is wrong). Service types (ClusterIP, NodePort, LoadBalancer, ExternalName) are part of the Service API and operate at Layer 4. Ingress is a separate API object that depends on an Ingress Controller to actually implement routing. The controller watches Ingress resources and configures a reverse proxy/load balancer (like NGINX, HAProxy, or a cloud load balancer integration) to enforce the desired routing. Without an Ingress Controller, creating an Ingress object alone will not route traffic.

Option A describes endpoint tracking (that's closer to Endpoints/EndpointSlice). Option C describes NetworkPolicy, which controls allowed network flows between Pods/namespaces. Ingress is about exposing and routing incoming application traffic from outside the cluster to internal Services.

So the verified correct statement is D: Ingress exposes routes from outside the cluster to Services in the cluster.

Question: 30

What best describes cloud native service discovery?

- A. It's a mechanism for applications and microservices to locate each other on a network.
- B. It's a procedure for discovering a MAC address, associated with a given IP address.
- C. It's used for automatically assigning IP addresses to devices connected to the network.
- D. It's a protocol that turns human-readable domain names into IP addresses on the Internet.

Answer: A

Explanation:

Cloud native service discovery is fundamentally about how services and microservices find and connect to each other reliably in a dynamic environment, so A is correct. In cloud native systems (especially Kubernetes), instances are ephemeral: Pods can be created, destroyed, rescheduled, and scaled at any time. Hardcoding IPs breaks quickly. Service discovery provides stable names and lookup mechanisms so that one component can locate another even as underlying endpoints change.

In Kubernetes, service discovery is commonly achieved through Services (stable virtual IP + DNS name) and cluster DNS (CoreDNS). A Service selects a group of Pods via labels, and Kubernetes maintains the set of endpoints behind that Service. Clients connect to the Service name (DNS) and Kubernetes routes traffic to the current healthy Pods. For some workloads, headless Services provide DNS records that map directly to Pod IPs for per-instance discovery.

The other options describe different networking concepts: B is ARP (MAC discovery), C is DHCP (IP assignment), and D is DNS in a general internet sense. DNS is often used as a mechanism for service discovery, but cloud native service discovery is broader: it's the overall mechanism enabling dynamic location of services, often implemented via DNS and/or environment variables and sometimes enhanced by service meshes.

So the best description remains A: a mechanism that allows applications and microservices to locate each other on a network in a dynamic environment.

Question: 31

What components are common in a service mesh?

- A. Tracing and log storage
- B. Circuit breaking and Pod scheduling
- C. Data plane and runtime plane
- D. Service proxy and control plane

Answer: D

Explanation:

A service mesh is an architectural pattern that manages service-to-service communication in a microservices environment by inserting a dedicated networking layer. The two most common building blocks you'll see across service mesh implementations are (1) a data plane of proxies and (2) a control plane that configures and manages those proxies—this aligns best with “service proxy and control plane,” option D.

In practice, the data plane is usually implemented via sidecar proxies (or sometimes node/ambient proxies) that sit “next to” workloads and handle traffic functions such as mTLS encryption, retries, timeouts, load balancing policies, traffic splitting, and telemetry generation. These proxies can capture inbound and outbound traffic without requiring changes to application code, which is one of the defining benefits of a mesh.

The control plane provides the management layer: it distributes policy and configuration to the proxies (routing rules, security policies, identities/certificates), discovers services/endpoints, and often coordinates certificate rotation and workload identity. In Kubernetes environments, meshes typically integrate with the Kubernetes API for service discovery and configuration.

Option C is close in spirit but uses non-standard wording (“runtime plane” is not a typical service mesh term; “control plane” is). Options A and B describe capabilities that may exist in a mesh ecosystem (telemetry, circuit breaking), but they are not the universal “core components” across meshes. Tracing/log storage, for example, is usually handled by external observability backends (e.g., Jaeger, Tempo, Loki) rather than being intrinsic “mesh components.”

So, the most correct and broadly accepted answer is D: service proxy and control plane.

Question: 32

Which storage operator in Kubernetes can help the system to self-scale, self-heal, etc?

- A. Rook
- B. Kubernetes
- C. Helm
- D. Container Storage Interface (CSI)

Answer: A

Explanation:

Rook is a Kubernetes storage operator that helps manage and automate storage systems in a Kubernetes-native way, so A is correct. The key phrase in the question is “storage operator ... selfscale, self-heal.” Operators extend Kubernetes by using controllers to reconcile a desired state. Rook applies that model to storage, commonly by managing storage backends like Ceph (and other systems depending on configuration).

With an operator approach, you declare how you want storage to look (cluster size, pools, replication, placement, failure domains), and the operator works continuously to maintain that state. That includes operational behaviors that feel “self-healing” such as reacting to failed storage Pods, rebalancing, or restoring desired replication counts (the exact behavior depends on the backend and configuration). The important KCNA-level idea is that Rook uses Kubernetes controllers to automate day-2 operations for storage in a way consistent with Kubernetes’ reconciliation loops.

The other options do not match the

question:
“Kubernetes

” is the orchestrator itself, not a storage operator. “Helm” is a package manager for Kubernetes apps—it can install storage software, but it is not an operator that continuously reconciles and selfmanages. “CSI” (Container Storage Interface) is an interface specification that enables pluggable storage drivers; CSI drivers provision and attach volumes, but CSI itself is not a “storage operator” with the broader self-managing operator semantics described here.

So, for “storage operator that can help with self-* behaviors,” Rook is the correct choice.

Question: 33

What fields must exist in any Kubernetes object (e.g. YAML) file?

- A. apiVersion, kind, metadata
- B. kind, namespace, data
- C. apiVersion, metadata, namespace
- D. kind, metadata, data

Answer: A

Explanation:

Any Kubernetes object manifest must include apiVersion, kind, and metadata, which makes A correct. This comes directly from how Kubernetes resources are represented and processed by the API server.

apiVersion tells Kubernetes which API group and version should be used to interpret the object (for example v1, apps/v1,

batch/v1). This matters because schemas and available fields can change between versions.

kind specifies the type of object you are creating (for example Pod, Service, Deployment, ConfigMap). Kubernetes uses this to route the request to the correct API endpoint and schema.

metadata contains identifying and organizational information such as name, namespace (when namespaced), labels, and annotations. At minimum, most objects require a name; labels and annotations are optional but extremely common for selection and tooling.

A common point of confusion is spec. Many Kubernetes objects include spec because they define desired state (like a Deployment's replica count, Pod template, update strategy). However, the question asks what fields must exist in any Kubernetes object file. Not all objects require a spec in the same way (and some objects include other top-level sections like data for ConfigMaps/Secrets or rules for RBAC objects). The truly universal top-level requirements are the trio in option A.

Options B, C, and D include fields that are not universally required (namespace is not required for cluster-scoped objects, and data only applies to certain kinds like ConfigMaps/Secrets). Therefore, apiVersion + kind + metadata is the correct, general rule and matches Kubernetes object structure.

Question: 34

Which of the following would fall under the responsibilities of an SRE?

- A. Developing a new application feature.
- B. Creating a monitoring baseline for an application.
- C. Submitting a budget for running an application in a cloud.
- D. Writing policy on how to submit a code change.

Answer: B

Explanation:

Site Reliability Engineering (SRE) focuses on reliability, availability, performance, and operational excellence using engineering approaches. Among the options, creating a monitoring baseline for an application is a classic SRE

responsibility, so B is correct. A monitoring baseline typically includes defining key service-level signals (latency, traffic, errors, saturation), establishing dashboards, setting sensible alert thresholds, and ensuring telemetry is complete enough to support incident response and capacity planning.

In Kubernetes environments, SRE work often involves ensuring that workloads expose health endpoints for probes, that resource requests/limits are set to allow stable scheduling and autoscaling, and that observability pipelines (metrics, logs, traces) are consistent. Building a monitoring baseline also ties into SLO/SLI practices: SREs define what “good” looks like, measure it continuously, and create alerts that notify teams when the system deviates from those expectations.

Option A is primarily an application developer task—SREs may contribute to reliability features, but core product feature development is usually owned by engineering teams. Option C is more aligned with finance, FinOps, or management responsibilities, though SRE data can inform costs. Option D is closer to governance, platform policy, or developer experience/process ownership; SREs might influence processes, but “policy on how to submit code change” is not the defining SRE duty compared to monitoring and reliability engineering.

Therefore, the best verified choice is B, because establishing monitoring baselines is central to operating reliable services on Kubernetes.

Question: 35

What are the initial namespaces that Kubernetes starts with?

- A. default, kube-system, kube-public, kube-node-lease
- B. default, system, kube-public
- C. kube-default, kube-system, kube-main, kube-node-lease
- D. kube-default, system, kube-main, kube-primary

Answer: A

Explanation:

Kubernetes creates a set of namespaces by default when a cluster is initialized. The standard initial namespaces are default, kube-system, kube-public, and kube-node-lease, making A correct.

default is the namespace where resources are created if you don't specify another namespace. Many quick-start examples deploy here, though production environments typically use dedicated namespaces per app/team.

kube-system contains objects created and managed by Kubernetes system components (control plane add-ons, system Pods, controllers, DNS components, etc.). It's a critical namespace, and access is typically restricted.

kube-public is readable by all users (including unauthenticated users in some configurations) and is intended for public cluster information, though it's used sparingly in many environments.

kube-node-lease holds Lease objects used for node heartbeats. This improves scalability by reducing

load on etcd compared to older heartbeat mechanisms and helps the control plane track node liveness efficiently.

The incorrect options contain non-standard namespace names like "system," "kube-main," or "kube- primary," and "kube-default" is not a real default namespace. Kubernetes' built-in namespace set is well-documented and consistent with typical cluster bootstraps.

Understanding these namespaces matters operationally: system workloads and controllers often live in kube-system, and many troubleshooting steps involve inspecting Pods and events there.

Meanwhile, kube-node-lease is key to node health tracking, and default is the catch-all if you forget to specify -n.

So, the verified answer is A: default, kube-system, kube-public, kube-node-lease.

Question: 36

What is a probe within Kubernetes?

- A. A monitoring mechanism of the Kubernetes API.
- B. A pre-operational scope issued by the kubectl agent.
- C. A diagnostic performed periodically by the kubelet on a container.
- D. A logging mechanism of the Kubernetes API.

Answer: C

Explanation:

In Kubernetes, a probe is a health check mechanism that the kubelet executes against containers, so C is correct. Probes are part of how Kubernetes implements self-healing and safe traffic management. The kubelet runs probes periodically according to the configuration in the Pod spec and uses the results to decide whether a container is healthy, ready to receive traffic, or still starting up.

Kubernetes supports three primary probe types:

Liveness probe: determines whether the container should be restarted. If liveness fails repeatedly, kubelet restarts the container (subject to restartPolicy).

Readiness probe: determines whether the Pod should receive traffic via Services. If readiness fails, the Pod is removed from Service endpoints, preventing traffic from being routed to it until it becomes ready again.

Startup probe: used for slow-starting containers. It disables liveness/readiness failures until startup succeeds, preventing premature restarts during initialization.

Probe mechanisms can be HTTP GET, TCP socket checks, or exec commands run inside the container. These checks are performed by kubelet on the node where the Pod is running, not by the API server.

Options A and D incorrectly attribute probes to the Kubernetes API. While probe configuration is stored in the API as part of Pod specs, execution is node-local. Option B is not a Kubernetes concept.

So the correct definition is: a probe is a periodic diagnostic run by kubelet to assess container health/readiness, enabling reliable rollouts, traffic gating, and automatic recovery.

Question: 37

Which Kubernetes feature would you use to guard against split brain scenarios with your distributed application?

- A. Replication controllers
- B. Consensus protocols
- C. Rolling updates
- D. StatefulSet

Answer: D

Explanation:

The exam-expected Kubernetes feature here is StatefulSet, so D is the correct answer. StatefulSets are designed for distributed/stateful applications that require stable network identities, stable storage, and ordered deployment/termination. Those properties are commonly required by systems that must avoid “split brain” behaviors—where multiple nodes believe they are the leader/primary due to partitions or identity confusion.

StatefulSets give each Pod a persistent identity (e.g., app-0, app-1) and stable DNS naming (typically via a headless Service), which supports consistent peer discovery and membership. They also commonly pair with PersistentVolumeClaims so that each replica keeps its own data across restarts and reschedules. The ordered rollout semantics help clustered systems bootstrap and expand in controlled sequences, reducing the chance of chaotic membership changes.

Important nuance: StatefulSet alone does not magically prevent split brain. Split brain prevention is primarily a property of the application’s own clustering/consensus design (e.g., leader election, quorum, fencing). That’s why option B (“consensus protocols”) is conceptually the true prevention mechanism—but it’s not a Kubernetes feature in the way the question frames it. Kubernetes provides primitives that make it feasible to run such systems safely (stable IDs, stable storage, predictable DNS), and StatefulSet is the Kubernetes workload API designed for that class of distributed stateful apps.

Replication controllers and rolling updates don’t address identity/quorum concerns. Therefore, within Kubernetes constructs, StatefulSet is the best verified choice for workloads needing stable identity patterns commonly used to reduce split-brain risk.

Question: 38

What feature must a CNI support to control specific traffic flows for workloads running in Kubernetes?

- A. Border Gateway Protocol
- B. IP Address Management
- C. Pod Security Policy
- D. Network Policies

Answer: D

Explanation:

To control which workloads can communicate with which other workloads in Kubernetes, you use NetworkPolicy resources—but enforcement depends on the cluster's networking implementation. Therefore, for traffic-flow control, the CNI/plugin must support Network Policies, making D correct.

Kubernetes defines the NetworkPolicy API as a declarative way to specify allowed ingress and egress traffic based on selectors (Pod labels, namespaces, IP blocks) and ports/protocols. However, Kubernetes itself does not enforce NetworkPolicy rules; enforcement is provided by the network plugin (or associated dataplane components). If your CNI does not implement NetworkPolicy, the objects may exist in the API but have no effect—Pods will communicate freely by default.

Option B (IP Address Management) is often part of CNI responsibilities, but IPAM is about assigning addresses, not enforcing L3/L4 security policy. Option A (BGP) is used by some CNIs to advertise routes (for example, in certain Calico deployments), but BGP is not the general requirement for policy enforcement. Option C (Pod Security Policy) is a deprecated/removed Kubernetes admission feature related to Pod security settings, not network flow control.

From a Kubernetes security standpoint, NetworkPolicies are a key tool for implementing least privilege at the network layer—limiting lateral movement, reducing blast radius, and segmenting environments. But they only work when the chosen CNI supports them. Thus, the correct answer is D: Network Policies.

Question: 39

What is the main role of the Kubernetes DNS within a cluster?

- A. Acts as a DNS server for virtual machines that are running outside the cluster.
- B. Provides a DNS as a Service, allowing users to create zones and registries for domains that they own.
- C. Allows Pods running in dual stack to convert IPv6 calls into IPv4 calls.
- D. Provides consistent DNS names for Pods and Services for workloads that need to communicate with each other.

Answer: D

Explanation:

Kubernetes DNS (commonly implemented by CoreDNS) provides service discovery inside the cluster by assigning stable, consistent DNS names to Services and (optionally) Pods, which makes D correct. In a Kubernetes environment, Pods are ephemeral—IP addresses can change when Pods restart or move between nodes. DNS-based discovery allows applications to communicate using stable names rather than hardcoded IPs.

For Services, Kubernetes creates DNS records like `service-name.namespace.svc.cluster.local`, which resolve to the Service's virtual IP (ClusterIP) or, for headless Services, to the set of Pod endpoints. This supports both load-balanced communication (standard Service) and per-Pod addressing (headless Service, commonly used with StatefulSets).

Kubernetes DNS is therefore a core building block that enables microservices to locate each other reliably.

Option A is not Kubernetes DNS's purpose; it serves cluster workloads rather than external VMs. Option B describes a managed DNS hosting product (creating zones/registries), which is outside the scope of cluster DNS. Option C describes protocol translation, which is not the role of DNS. Dual-stack support relates to IP families and networking configuration, not DNS translating IPv6 to IPv4.

In day-to-day Kubernetes operations, DNS reliability impacts everything: if DNS is unhealthy, Pods may fail to resolve Services, causing cascading outages. That's why CoreDNS is typically deployed as a highly available add-on in kube-system, and why DNS caching and scaling are important for large clusters.

So the correct statement is D: Kubernetes DNS provides consistent DNS names so workloads can communicate reliably.

Question: 40

Scenario: You have a Kubernetes cluster hosted in a public cloud provider. When trying to create a Service of type LoadBalancer, the external-ip is stuck in the "Pending" state. Which Kubernetes component is failing in this scenario?

- A. Cloud Controller Manager
- B. Load Balancer Manager
- C. Cloud Architecture Manager
- D. Cloud Load Balancer Manager

Answer: A

Explanation:

When you create a Service of type LoadBalancer in a cloud environment, Kubernetes relies on cloudprovider integration to provision an external load balancer and allocate a public IP (or equivalent). The control plane component responsible for this integration is the cloud-controller-manager, so A is CORRECT.

In Kubernetes, a LoadBalancer Service triggers a controller loop that calls the cloud provider APIs to create/update a load balancer that forwards traffic to the cluster (often via NodePorts on worker nodes, or via provider-specific mechanisms). The Service remains with EXTERNAL-IP: Pending until the cloud provider resource is successfully created and the controller updates the Service status with the assigned external address. If that status never updates, it usually indicates the cloud integration path is broken—commonly due to: missing cloud provider configuration, broken credentials/IAM permissions, the cloud-controller-manager not running/healthy, or a misconfigured cloud provider implementation.

The other options are not real Kubernetes components. Kubernetes does not include a “Load Balancer Manager” or “Cloud Architecture Manager” component name in its standard architecture. In many managed Kubernetes offerings, the cloud-controller-manager (or its equivalent) is provided/managed by the provider, but the responsibility remains the same: reconcile Kubernetes Service resources into cloud load balancer resources.

Therefore, in this scenario, the failing component is the Cloud Controller Manager, which is the Kubernetes control plane component that interfaces with the cloud provider to provision external load balancers and update the Service status.

Question: 41

What are the characteristics for building every cloud-native application?

- A. Resiliency, Operability, Observability, Availability
- B. Resiliency, Containerd, Observability, Agility
- C. Kubernetes, Operability, Observability, Availability
- D. Resiliency, Agility, Operability, Observability

Answer: D

Explanation:

Cloud-native applications are typically designed to thrive in dynamic, distributed environments where infrastructure is elastic and failures are expected. The best set of characteristics listed is Resiliency, Agility, Operability,

Observability, making D correct.

Resiliency means the application and its supporting platform can tolerate failures and continue providing service. In Kubernetes terms, resiliency is supported through self-healing controllers, replica management, health probes, and safe rollout mechanisms, but the application must also be designed to handle transient failures, retries, and graceful degradation.

Agility reflects the ability to deliver changes quickly and safely. Cloud-native systems emphasize automation, CI/CD, declarative configuration, and small, frequent releases—often enabled by Kubernetes primitives like Deployments and rollout strategies. Agility is about reducing the friction to ship improvements while maintaining reliability.

Operability is how manageable the system is in production: clear configuration, predictable deployments, safe scaling, and automation-friendly operations. Kubernetes encourages operability through consistent APIs, controllers, and standardized patterns for configuration and lifecycle.

Observability means you can understand what's happening inside the system using telemetry— metrics, logs, and traces—so you can troubleshoot issues, measure SLOs, and improve performance. Kubernetes provides many integration points for observability, but cloud-native apps must also emit meaningful signals.

Options B and C include items that are not “characteristics” (containerd is a runtime; Kubernetes is a platform). Option A includes “availability,” which is important, but the canonical cloud-native framing in this question emphasizes the four qualities in D as the foundational build characteristics.

Question: 42

What does CNCF stand for?

- A. Cloud Native Community Foundation
- B. Cloud Native Computing Foundation
- C. Cloud Neutral Computing Foundation
- D. Cloud Neutral Community Foundation

Answer: B

Explanation:

CNCF stands for the Cloud Native Computing Foundation, making B correct. CNCF is the foundation that hosts and sustains many cloud-native open source projects, including Kubernetes, and provides governance, neutral stewardship,

and community infrastructure to help projects grow and remain vendor-neutral.

CNCF's scope includes not only Kubernetes but also a broad ecosystem of projects across observability, networking, service meshes, runtime security, CI/CD, and application delivery. The foundation defines processes for project incubation and graduation, promotes best practices, organizes community events, and supports interoperability and adoption through reference architectures and education.

In the Kubernetes context, CNCF's role matters because Kubernetes is a massive multi-vendor project. Neutral governance reduces the risk that any single company can unilaterally control direction. This fosters broad contribution and adoption across cloud providers and enterprises. CNCF also supports the broader "cloud native" definition, often associated with containerization, microservices, declarative APIs, automation, and resilience principles.

The incorrect options are close-sounding but not accurate expansions. "Cloud Native Community Foundation" and the "Cloud Neutral ..." variants are not the recognized meaning. The correct official name is Cloud Native Computing Foundation.

So, the verified answer is B, and understanding CNCF helps connect Kubernetes to its broader ecosystem of standardized, interoperable cloud-native tooling.

Question: 43

Kubernetes supports multiple virtual clusters backed by the same physical cluster. These virtual clusters are called:

- A. Namespaces
- B. Containers
- C. Hypervisors
- D. cgroups

Answer: A

Explanation:

Kubernetes provides "virtual clusters" within a single physical cluster primarily through Namespaces, so A is correct.

Namespaces are a logical partitioning mechanism that scopes many Kubernetes resources (Pods, Services, Deployments, ConfigMaps, Secrets, etc.) into separate environments. This enables multiple teams, applications, or environments (dev/test/prod) to share a cluster while keeping their resource names and access controls separated.

Namespaces are often described as “soft multi-tenancy.” They don’t provide full isolation like separate clusters, but they do allow administrators to apply controls per namespace:

RBAC rules can grant different permissions per namespace (who can read Secrets, who can deploy workloads, etc.).

ResourceQuotas and LimitRanges can enforce fair usage and prevent one namespace from consuming all cluster resources.

NetworkPolicies can isolate traffic between namespaces (depending on the CNI).

Containers are runtime units inside Pods and are not “virtual clusters.” Hypervisors are virtualization components for VMs, not Kubernetes partitioning constructs. cgroups are Linux kernel primitives for resource control, not Kubernetes virtual cluster constructs.

While there are other “virtual cluster” approaches (like vcluster projects) that create stronger virtualized control planes, the built-in Kubernetes mechanism referenced by this question is namespaces. Therefore, the correct answer is

A: Namespaces.

Question: 44

What component enables end users, different parts of the Kubernetes cluster, and external components to communicate with one another?

- A. kubectl
- B. AWS Management Console
- C. Kubernetes API
- D. Google Cloud SDK

Answer: C

Explanation:

The Kubernetes API is the central interface that enables communication between users, controllers, nodes, and external integrations, so C is correct. Kubernetes is fundamentally an API-driven system: all cluster state is represented as API objects, and all operations—create, update, delete, watch—flow through the API server.

End users typically interact with the Kubernetes API using tools like kubectl, client libraries, or dashboards. But those tools are clients; the shared communication “hub” is the API itself. Inside the cluster, core control plane components (controllers, scheduler) continuously watch the API for desired state and write status updates back. Worker nodes (via kubelet) also communicate with the API server to receive Pod specs, report node health, and update Pod statuses.

External systems— cloud provider integrations, CI/CD pipelines, GitOps controllers, monitoring and policy engines—also integrate primarily through the Kubernetes API.

Option A (kubectl) is a CLI that talks to the Kubernetes API; it is not the underlying component that all parts use to communicate. Options B and D are cloud-provider tools and are not universal to Kubernetes clusters. Kubernetes runs across many environments, and the consistent interoperability layer is the Kubernetes API.

This API-centric architecture is what enables Kubernetes’ declarative model: you submit desired state to the API, and controllers reconcile actual state to match. It also enables extensibility: CRDs and admission webhooks expand what the API can represent and enforce. Therefore, the correct answer is C: Kubernetes API.

Question: 45

Which command will list the resource types that exist within a cluster?

- A. `kubectl api-resources`
- B. `kubectl get namespaces`
- C. `kubectl api-versions`
- D. `curl https://kubectrl/namespaces`

Answer: A

Explanation:

To list the resource types available in a Kubernetes cluster, you use `kubectl api-resources`, so A is correct. This command queries the API server’s discovery endpoints and prints a table of resources (kinds) that the cluster knows about, including

their names, shortnames, API group/version, whether they are namespaced, and supported verbs. It's extremely useful for learning what objects exist in a cluster—especially when CRDs are installed, because those custom resource types will also appear in the output.

Option C (`kubectl api-versions`) lists available API versions (group/version strings like `v1`, `apps/v1`,

`batch/v1`) but does not directly list the resource kinds/types. It's related discovery information but answers a different question. Option B (`kubectl get namespaces`) lists namespaces, not resource types. Option D is invalid (typo in URL and conceptually not the Kubernetes discovery mechanism).

Practically, `kubectl api-resources` is used during troubleshooting and exploration: you might use it to confirm whether a CRD is installed (e.g., `certificates.cert-manager.io` kinds), to check whether a resource is namespaced, or to find the correct kind name for `kubectl get`. It also helps understand what your cluster supports at the API layer (including aggregated APIs).

So, the verified correct command to list resource types that exist in the cluster is A: `kubectl api-resources`.

Question: 46

Which of these components is part of the Kubernetes Control Plane?

- A. CoreDNS
- B. cloud-controller-manager
- C. kube-proxy
- D. kubelet

Answer: B

Explanation:

The Kubernetes control plane is the set of components responsible for making cluster-wide decisions (like scheduling) and detecting and responding to cluster events (like starting new Pods when they fail). In upstream Kubernetes architecture, the canonical control plane components include `kube-apiserver`, `etcd`, `kube-scheduler`, and `kube-controller-manager`, and—when running on a cloud provider—the `cloud-controller-manager`. That makes option B the correct answer: `cloud-controller-manager` is explicitly a control plane component that integrates Kubernetes with the underlying

cloud.

The cloud-controller-manager runs controllers that talk to cloud APIs for infrastructure concerns such as node lifecycle, routes, and load balancers. For example, when you create a Service of type

LoadBalancer, a controller in this component is responsible for provisioning a cloud load balancer and updating the Service status. This is clearly control-plane behavior: reconciling desired state into real infrastructure state.

Why the others are not control plane components (in the classic classification): kubelet is a node component (agent) responsible for running and managing Pods on a specific node. kube-proxy is also a node component that implements Service networking rules on nodes. CoreDNS is usually deployed as a cluster add-on for DNS-based service discovery; it's critical, but it's not a control plane component in the strict architectural list.

So, while many clusters run CoreDNS in kube-system, the Kubernetes component that is definitively "part of the control plane" among these choices is cloud-controller-manager (B).

Question: 47

Which of the following systems is NOT compatible with the CRI runtime interface standard?

(Typo corrected: "CRI-0" → "CRI-O")

- A. CRI-O
- B. dockershim
- C. systemd
- D. containerd

Answer: C

Explanation:

Kubernetes uses the Container Runtime Interface (CRI) to support pluggable container runtimes. The kubelet talks to a CRI-compatible runtime via gRPC, and that runtime is responsible for pulling images and running containers. In this context, containerd and CRI-O are CRI-compatible container runtimes (or runtime stacks) used widely with Kubernetes,

and dockershim historically served as a compatibility layer that allowed kubelet to talk to Docker Engine as if it were CRI (before dockershim was removed from kubelet in newer Kubernetes versions). That leaves systemd as the correct "NOT compatible with CRI" answer, so C is correct.

systemd is an init system and service manager for Linux. While it can be involved in how services (like kubelet) are started and managed on the host, it is not a container runtime implementing CRI. It does not provide CRI gRPC endpoints for kubelet, nor does it manage containers in the CRI sense.

The deeper Kubernetes concept here is separation of responsibilities: kubelet is responsible for Pod lifecycle at the node level, but it delegates "run containers" to a runtime via CRI. Runtimes like containerd and CRI-O implement that contract; Kubernetes can swap them without changing kubelet logic. Historically, dockershim translated kubelet's CRI calls into Docker Engine calls. Even though dockershim is no longer part of kubelet, it was still "CRI-adjacent" in purpose and often treated as compatible in older curricula.

Therefore, among the provided options, systemd is the only one that is clearly not a CRI-compatible runtime system, making C correct.

Question: 48

What is a key feature of a container network?

- A. Proxying REST requests across a set of containers.
- B. Allowing containers running on separate hosts to communicate.
- C. Allowing containers on the same host to communicate.
- D. Caching remote disk access.

Answer: B

Explanation:

A defining requirement of container networking in orchestrated environments is enabling workloads to communicate across hosts, not just within a single machine. That's why B is correct: a key feature of a container network is allowing

containers (Pods) running on separate hosts to communicate.

In Kubernetes, this idea becomes the Kubernetes network model: every Pod gets an IP address, and Pods should be able to communicate with other Pods across nodes without needing NAT (depending on implementation details). Achieving that across a cluster requires a networking layer (typically implemented by a CNI plugin) that can route traffic between nodes so that Pod-to-Pod communication works regardless of placement. This is crucial because schedulers dynamically place Pods; you cannot assume two communicating components will land on the same node.

Option C is true in a trivial sense—containers on the same host can communicate—but that capability alone is not the key feature that makes orchestration viable at scale. Cross-host connectivity is the harder and more essential property.

Option A describes application-layer behavior (like API gateways or reverse proxies) rather than the foundational networking capability. Option D describes storage optimization, unrelated to container networking.

From a cloud native architecture perspective, reliable cross-host networking enables microservices patterns, service discovery, and distributed systems behavior. Kubernetes Services, DNS, and NetworkPolicies all depend on the underlying ability for Pods across the cluster to send traffic to each other. If your container network cannot provide cross-node routing and reachability, the cluster behaves like isolated islands and breaks the fundamental promise of orchestration:

“schedule anywhere, communicate consistently.”

Question: 49

How can you monitor the progress for an updated Deployment/DaemonSets/StatefulSets?

- A. kubectl rollout watch
- B. kubectl rollout progress
- C. kubectl rollout state
- D. kubectl rollout status

Answer: D

Explanation:

To monitor rollout progress for Kubernetes workload updates (most commonly Deployments, and also StatefulSets and DaemonSets where applicable), the standard kubectl command is kubectl rollout status, which makes D correct.

Kubernetes manages updates declaratively through controllers. For a Deployment, an update typically creates a new ReplicaSet and gradually shifts replicas from the old to the new according to the strategy (e.g., RollingUpdate with maxUnavailable and maxSurge). For StatefulSets, updates may be ordered and respect stable identities, and for DaemonSets, an update replaces node-level Pods according to update strategy. In all cases, you often want a single command that tells you whether the controller has completed the update and whether the new replicas are available. `kubectl rollout status` queries the resource status and prints a progress view until completion or timeout.

The other commands listed are not the canonical `kubectl` subcommands. `kubectl rollout watch`, `kubectl rollout progress`, and `kubectl rollout state` are not standard rollout verbs in `kubectl`. The supported rollout verbs typically include `status`, `history`, `undo`, `pause`, and `resume` (depending on `kubectl` version and resource type).

Operationally, `kubectl rollout status deployment/<name>` is used during releases and troubleshooting to confirm that new Pods become ready, that old replicas scale down, and that the rollout does not stall due to readiness probe failures, insufficient resources, or policy constraints (e.g., `PodDisruptionBudgets` interacting with drains). It ties directly into Kubernetes' self-healing and declarative update model: you declare a new desired Pod template, and controllers reconcile toward it while `rollout status` reports that reconciliation's progress.

Question: 50

What is the goal of load balancing?

- A. Automatically measure request performance across instances of an application.
- B. Automatically distribute requests across different versions of an application.
- C. Automatically distribute instances of an application across the cluster.
- D. Automatically distribute requests across instances of an application.

Answer: D

Explanation:

The core goal of load balancing is to distribute incoming requests across multiple instances of a service so that no single instance becomes overloaded and so that the overall service is more available and responsive. That matches option D, which is the correct answer.

In Kubernetes, load balancing commonly appears through the Service abstraction. A Service selects a set of Pods using

labels and provides stable access via a virtual IP (ClusterIP) and DNS name. Traffic sent to the Service is then forwarded to one of the healthy backend Pods. This spreads load across replicas and provides resilience: if one Pod fails, it is removed from endpoints (or becomes NotReady) and traffic shifts to remaining replicas. The actual traffic distribution mechanism depends on the networking implementation (kube-proxy using iptables/IPVS or an eBPF dataplane), but the intent remains consistent: distribute requests across multiple backends.

Option A describes monitoring/observability, not load balancing. Option B describes progressive delivery patterns like canary or A/B routing; that can be implemented with advanced routing layers (Ingress controllers, service meshes), but it's not the general definition of load balancing. Option C describes scheduling/placement of instances (Pods) across cluster nodes, which is the role of the scheduler and controllers, not load balancing.

In cloud environments, load balancing may also be implemented by external load balancers (cloud LBs) in front of the cluster, then forwarded to NodePorts or ingress endpoints, and again balanced internally to Pods. At each layer, the objective is the same: spread request traffic across multiple service instances to improve performance and availability.

Question: 51

How are ReplicaSets and Deployments related?

- A. Deployments manage ReplicaSets and provide declarative updates to Pods.
- B. ReplicaSets manage stateful applications, Deployments manage stateless applications.
- C. Deployments are runtime instances of ReplicaSets.
- D. ReplicaSets are subsets of Jobs and CronJobs which use imperative Deployments.

Answer: A

Explanation:

In Kubernetes, a Deployment is a higher-level controller that manages ReplicaSets, and ReplicaSets in turn manage Pods. That is exactly what option A states, making it the correct answer.

A ReplicaSet's job is straightforward: ensure that a specified number of Pod replicas matching a selector are running. It continuously reconciles actual state to desired state by creating new Pods when replicas are missing or removing Pods when there are too many. However, ReplicaSets alone do not provide the richer application rollout lifecycle features

most teams need.

A Deployment adds those features by managing ReplicaSets across versions of your Pod template. When you update a Deployment (for example, change the container image tag), Kubernetes creates a new ReplicaSet with the new Pod template and then gradually scales the new ReplicaSet up and the old one down according to the Deployment strategy (RollingUpdate by default). Deployments also maintain rollout history, support rollback (kubectl rollout undo), and allow pause/resume of rollouts. This is why the common guidance is: you almost always create Deployments rather than ReplicaSets directly for stateless apps.

Option B is incorrect because stateful workloads are typically handled by StatefulSets, not ReplicaSets. Deployments can run stateless apps, but ReplicaSets are also used under Deployments and are not “for stateful only.” Option C is reversed: ReplicaSets are not “instances” of Deployments; Deployments create/manage ReplicaSets. Option D is incorrect because Jobs/CronJobs are separate controllers for run-to-completion workloads and do not define ReplicaSets as subsets.

So the accurate relationship is: Deployment → manages ReplicaSets → which manage Pods, enabling declarative updates and controlled rollouts.

Question: 52

What factors influence the Kubernetes scheduler when it places Pods on nodes?

- A. Pod memory requests, node taints, and Pod affinity.
- B. Pod labels, node labels, and request labels.
- C. Node taints, node level, and Pod priority.
- D. Pod priority, container command, and node labels.

Answer: A

Explanation:

The Kubernetes scheduler chooses a node for a Pod by evaluating scheduling constraints and cluster state. Key inputs include resource requests (CPU/memory), taints/tolerations, and affinity/anti-affinity rules. Option A directly names three real, high-impact scheduling factors—Pod memory requests, node taints, and Pod affinity—so A is

correct.

Resource requests are fundamental: the scheduler must ensure the target node has enough allocatable CPU/memory to satisfy the Pod's requests. Requests (not limits) drive placement decisions. Taints on nodes repel Pods unless the Pod has a matching toleration, which is commonly used to reserve nodes for special workloads (GPU nodes, system nodes, restricted nodes) or to protect nodes under certain conditions. Affinity and anti-affinity allow expressing "place me near" or "place me away" rules—e.g., keep replicas spread across failure domains or co-locate components for latency.

Option B includes labels, which do matter, but "request labels" is not a standard scheduler concept; labels influence scheduling mainly through selectors and affinity, not as a direct category called "request labels." Option C mixes a real concept (taints, priority) with "node level," which isn't a standard scheduling factor term. Option D includes "container command," which does not influence scheduling; the scheduler does not care what command the container runs, only placement constraints and resources.

Under the hood, kube-scheduler uses a two-phase process (filtering then scoring) to select a node, but the inputs it filters/scores include exactly the kinds of constraints in A. Therefore, the verified best answer is A.

Question: 53

What is the core metric type in Prometheus used to represent a single numerical value that can go up and down?

- A. Summary
- B. Counter
- C. Histogram
- D. Gauge

Answer: D

Explanation:

In Prometheus, a Gauge represents a single numerical value that can increase and decrease over time, which makes D the correct answer. Gauges are used for values like current memory usage, number of in-flight requests, queue depth, temperature, or CPU usage—anything that can move up and down.

This contrasts with a Counter, which is strictly monotonically increasing (it only goes up, except for resets when a process restarts). Counters are ideal for cumulative totals like total HTTP requests served, total errors, or bytes transmitted.

Histograms and Summaries are used to capture distributions (often latency distributions), providing bucketed counts (histogram) or quantile approximations (summary), and are not the “single value that goes up and down” primitive the question asks for.

In Kubernetes observability, metrics are a primary signal for understanding system health and performance. Prometheus is widely used to scrape metrics from Kubernetes components (kubelet, API server, controller-manager), cluster add-ons, and applications. Gauges are common for resource utilization metrics and for instantaneous states, such as `container_memory_working_set_bytes` or `go_goroutines`.

When you build alerting and dashboards, selecting the right metric type matters. For example, if you want to alert on the current memory usage, a gauge is appropriate. If you want to compute request rates, you typically use counters with Prometheus functions like `rate()` to derive per-second rates. Histograms and summaries are used when you need latency percentiles or distribution analysis.

So, for “a single numerical value that can go up and down,” the correct Prometheus metric type is Gauge (D).

Question: 54

What is the primary mechanism to identify grouped objects in a Kubernetes cluster?

- A. Custom Resources
- B. Labels
- C. Label Selector
- D. Pod

Answer: B

Explanation:

Kubernetes groups and organizes objects primarily using labels, so B is correct. Labels are key-value pairs attached to objects (Pods, Deployments, Services, Nodes, etc.) and are intended to be used for identifying, selecting, and grouping resources in a flexible, user-defined way.

Labels enable many core Kubernetes behaviors. For example, a Service selects the Pods that should receive traffic by matching a label selector against Pod labels. A Deployment’s ReplicaSet similarly uses label selectors to determine which Pods belong to the replica set. Operators and platform tooling also rely on labels to group resources by application, environment, team, or cost center. This is why labeling is considered foundational Kubernetes hygiene: consistent labels

make automation, troubleshooting, and governance easier.

A “label selector” (option C) is how you query/group objects based on labels, but the underlying primary mechanism is still the labels themselves. Without labels applied to objects, selectors have nothing to match. Custom Resources (option A) extend the API with new kinds, but they are not the primary grouping mechanism across the cluster. “Pod” (option D) is a workload unit, not a grouping mechanism.

Practically, Kubernetes recommends common label keys like `app.kubernetes.io/name`, `app.kubernetes.io/instance`, and `app.kubernetes.io/part-of` to standardize grouping. Those conventions improve interoperability with dashboards, GitOps tooling, and policy engines.

So, when the question asks for the primary mechanism used to identify grouped objects in

Kubernetes, the most accurate answer is Labels (B)—they are the universal metadata primitive used to group and select resources.

Question: 55

What is the name of the Kubernetes resource used to expose an application?

- A. Port
- B. Service
- C. DNS
- D. Deployment

Answer: B

Explanation:

To expose an application running on Pods so that other components can reliably reach it, Kubernetes uses a Service, making B the correct answer. Pods are ephemeral: they can be recreated, rescheduled, and scaled, which means Pod IPs change. A Service provides a stable endpoint (virtual IP and DNS name) and load-balances traffic across the set of Pods selected by its label selector.

Services come in multiple forms. The default is ClusterIP, which exposes the application inside the cluster. NodePort exposes the Service on a static port on each node, and LoadBalancer (in supported clouds) provisions an external load balancer that routes traffic to the Service. ExternalName maps a Service name to an external DNS name. But across these variants, the abstraction is consistent: a Service defines how to access a logical group of Pods.

Option A (Port) is not a Kubernetes resource type; ports are fields within resources. Option C (DNS) is a supporting mechanism (CoreDNS creates DNS entries for Services), but DNS is not the resource you create to expose the app. Option D (Deployment) manages Pod replicas and rollouts, but it does not directly provide stable networking access; you typically pair a Deployment with a Service to expose it.

This is a core cloud-native pattern: controllers manage compute, Services manage stable connectivity, and higher-level gateways like Ingress provide L7 routing for HTTP/HTTPS. So, the Kubernetes resource used to expose an application is Service (B).

Question: 56

What is a DaemonSet?

- A. It's a type of workload that ensures a specific set of nodes run a copy of a Pod.
- B. It's a type of workload responsible for maintaining a stable set of replica Pods running in any node.
- C. It's a type of workload that needs to be run periodically on a given schedule.
- D. It's a type of workload that provides guarantees about ordering, uniqueness, and identity of a set of Pods.

Answer: A

Explanation:

A DaemonSet ensures that a copy of a Pod runs on each node (or a selected subset of nodes), which matches option A and makes it correct. DaemonSets are ideal for node-level agents that should exist everywhere, such as log shippers, monitoring agents, CNI components, storage daemons, and security scanners.

DaemonSets differ from Deployments/ReplicaSets because their goal is not "N replicas anywhere," but "one replica per node" (subject to node selection). When nodes are added to the cluster, the DaemonSet controller automatically schedules the DaemonSet Pod onto the new nodes. When nodes are removed, the Pods associated with those nodes are cleaned up. You can restrict placement using node selectors, affinity rules, or tolerations so that only certain nodes run

the DaemonSet (for example, only Linux nodes, only GPU nodes, or only nodes with a dedicated label).

Option B sounds like a ReplicaSet/Deployment behavior (stable set of replicas), not a DaemonSet. Option C describes CronJobs (scheduled, recurring run-to-completion workloads). Option D describes StatefulSets, which provide stable identity, ordering, and uniqueness guarantees for stateful replicas.

Operationally, DaemonSets matter because they often run critical cluster services. During maintenance and upgrades, DaemonSet update strategy determines how those node agents roll out across the fleet. Since DaemonSets can tolerate taints (like master/control-plane node taints), they can also be used to ensure essential agents run across all nodes, including special pools. Thus, the correct definition is A.

Question: 57

What is the telemetry component that represents a series of related distributed events that encode the end-to-end request flow through a distributed system?

- A. Metrics
- B. Logs
- C. Spans
- D. Traces

Answer: D

Explanation:

In observability, traces represent an end-to-end view of a request as it flows through multiple services, so D is correct.

Tracing is particularly important in cloud-native microservices architectures because a single user action (like “checkout” or “search”) may traverse many services via HTTP/gRPC calls, message queues, and databases. Traces link those related events together so you can see where time is spent, where errors occur, and how dependencies behave.

A trace is typically composed of multiple spans (option C). A span is a single timed operation (e.g., “HTTP GET /orders”, “DB query”, “call payment service”). Spans include timing, attributes (tags), status/error information, and parent/child relationships. While spans are essential building blocks, the “series of related distributed events encoding end-to-end request flow” is the trace as a whole, not an individual span.

Metrics (option A) are numeric time series used for aggregation and alerting (rates, latency percentiles when derived, resource usage). Logs (option B) are discrete event records (text or structured) useful for forensic detail and debugging. Both are valuable, but neither inherently provides a stitched, causal, end-to-end request path across services. Traces do exactly that by propagating trace context (trace IDs/span IDs) across service boundaries (often via headers).

In Kubernetes environments, traces are commonly exported via OpenTelemetry instrumentation/collectors and visualized in tracing backends. Tracing enables faster incident resolution by pinpointing the slow hop, the failing downstream dependency, or unexpected fan-out. Therefore, the correct telemetry component for end-to-end distributed request flow is Traces (D).

Question: 58

In the Kubernetes platform, which component is responsible for running containers?

- A. etcd
- B. CRI-O
- C. cloud-controller-manager
- D. kube-controller-manager

Answer: B

Explanation:

In Kubernetes, the actual act of running containers on a node is performed by the container runtime. The kubelet instructs the runtime via CRI, and the runtime pulls images, creates containers, and manages their lifecycle. Among the options provided, CRI-O is the only container runtime, so B is CORRECT.

It's important to be precise: the component that "runs containers" is not the control plane and not etcd. etcd (option A) stores cluster state (API objects) as the backing datastore. It never runs containers. cloud-controller-manager (option C) integrates with cloud APIs for infrastructure like load balancers and nodes. kube-controller-manager (option D) runs controllers that reconcile Kubernetes objects (Deployments, Jobs, Nodes, etc.) but does not execute containers on worker nodes.

CRI-O is a CRI implementation that is optimized for Kubernetes and typically uses an OCI runtime (like runc) under the hood to start containers. Another widely used runtime is containerd. The runtime is installed on nodes and is a prerequisite for kubelet to start Pods. When a Pod is scheduled to a node, kubelet reads the PodSpec and asks the runtime to create a “pod sandbox” and then start the container processes. Runtime behavior also includes pulling images, setting up namespaces/cgroups, and exposing logs/stdout streams back to Kubernetes tooling.

So while “the container runtime” is the most general answer, the question’s option list makes CRI-O the correct selection because it is a container runtime responsible for running containers in Kubernetes.

Question: 59

Services and Pods in Kubernetes are _____ objects.

- A. JSON
- B. YAML
- C. Java
- D. REST

Answer: D

Explanation:

In Kubernetes, resources like Pods and Services are represented as API objects that you create, read, update, delete, and watch via the Kubernetes RESTful API. That makes D (REST) the correct answer.

Kubernetes is fundamentally API-driven: the API server exposes endpoints for each resource type (for example, `/api/v1/namespaces/{ns}/pods` and `/api/v1/namespaces/{ns}/services`). Clients such as `kubectl`, controllers, operators, and external systems interact with these resources by making REST-style calls using HTTP verbs (GET, POST, PUT/PATCH, DELETE) and using watch streams for event-driven updates. This API-first design is what enables Kubernetes’ declarative model—users submit desired state to the API server, and controllers reconcile the cluster to that desired state.

Options A and B (JSON and YAML) are common serialization formats used to represent Kubernetes objects, but they are

not what the objects “are.” Kubernetes objects are logical API resources; they can be encoded as JSON (what the API uses) and often authored as YAML for human convenience.

YAML is effectively a superset-friendly format that can be converted to JSON. The underlying API object model remains the same regardless of whether you wrote YAML or JSON. Option C (Java) is unrelated; Java is a programming language that can interact with Kubernetes via client libraries, but Kubernetes objects are not “Java objects” in the platform’s definition.

So the accurate statement is: Pods and Services are Kubernetes REST API objects (resources) exposed and managed through the Kubernetes API server, which is why REST is the correct fill-in.

Question: 60

What Kubernetes component handles network communications inside and outside of a cluster, using operating system packet filtering if available?

- A. kube-proxy
- B. kubelet
- C. etcd
- D. kube-controller-manager

Answer: A

Explanation:

kube-proxy is the Kubernetes component responsible for implementing Service networking on nodes, commonly by programming operating system packet filtering / forwarding rules (like iptables or IPVS), which makes A correct.

Kubernetes Services provide stable virtual IPs and ports that route traffic to a dynamic set of Pod endpoints. kube-proxy watches the API server for Service and EndpointSlice/Endpoints updates and then configures the node’s networking so that traffic to a Service is correctly forwarded to one of the backend Pods. In iptables mode, kube-proxy installs NAT and forwarding rules; in IPVS mode, it

programs kernel load-balancing tables. In both cases, it leverages OS-level packet handling to efficiently steer traffic. This

is the “packet filtering if available” concept referenced in the question.

kube-proxy’s work affects both “inside” and “outside” paths in typical setups. Internal cluster clients reach Services via ClusterIP and DNS, and kube-proxy rules forward that traffic to Pods. For external traffic, paths often involve NodePort or LoadBalancer Services or Ingress controllers that ultimately forward into Services/Pods—again relying on node-level service rules. While some modern CNI/eBPF dataplanes can replace or bypass kube-proxy, the classic Kubernetes architecture still defines kube- proxy as the component implementing Service connectivity.

The other options are not networking dataplane components: kubelet runs Pods and reports status; etcd stores cluster state; kube-controller-manager runs control loops for API objects. None of these handle node-level packet routing for Services. Therefore, the correct verified answer is A: kube- proxy.

Question: 61

What Kubernetes control plane component exposes the programmatic interface used to create, manage and interact with the Kubernetes objects?

- A. kube-controller-manager
- B. kube-proxy
- C. kube-apiserver
- D. etcd

Answer: C

Explanation:

The kube-apiserver is the front door of the Kubernetes control plane and exposes the programmatic interface used to create, read, update, delete, and watch Kubernetes objects—so C is correct. Every interaction with cluster state ultimately goes through the Kubernetes API. Tools like kubectl, client libraries, GitOps controllers, operators, and core control plane components (scheduler and controllers) all communicate with the API server to submit desired state and to observe current state.

The API server is responsible for handling authentication (who are you?), authorization (what are you allowed to do?), and admission control (should this request be allowed and possibly mutated/validated?). After a request passes these

gates, the API server persists the object's desired state to etcd (the backing datastore) and returns a response. The API server also provides a watch mechanism so controllers can react to changes efficiently, enabling Kubernetes' reconciliation model.

It's important to distinguish this from the other options. etcd stores cluster data but does not expose the cluster's primary user-facing API; it's an internal datastore. kube-controller-manager runs control loops (controllers) that continuously reconcile resources (like Deployments, Nodes, Jobs) but it consumes the API rather than exposing it. kube-proxy is a node-level component implementing Service networking rules and is unrelated to the control-plane API endpoint.

Because Kubernetes is "API-driven," the kube-apiserver is central: if it is unavailable, you cannot create workloads, update configurations, or even reliably observe cluster state. This is why high availability architectures prioritize multiple API server instances behind a load balancer, and why securing the API server (RBAC, TLS, audit) is a primary operational concern.

Question: 62

Which type of Service requires manual creation of Endpoints?

- A. LoadBalancer
- B. Services without selectors
- C. NodePort
- D. ClusterIP with selectors

Answer: B

Explanation:

A Kubernetes Service without selectors requires you to manage its backend endpoints manually, so B

is correct. Normally, a Service uses a selector to match a set of Pods (by labels). Kubernetes then automatically maintains the backend list (historically Endpoints, now commonly EndpointSlice) by tracking which Pods match the selector and are Ready. This automation is one of the key reasons Services provide stable connectivity to dynamic Pods.

When you create a Service without a selector, Kubernetes has no way to know which Pods (or external IPs) should receive traffic. In that pattern, you explicitly create an Endpoints object (or EndpointSlices, depending on your approach and controller support) that maps the Service name to one or more IP:port tuples. This is commonly used to represent external services (e.g., a database running outside the cluster) while still providing a stable Kubernetes Service DNS name for in-cluster clients. Another use case is advanced migration scenarios where endpoints are controlled by custom controllers rather than label selection.

Why the other options are wrong: Service types like ClusterIP, NodePort, and LoadBalancer describe how a Service is exposed, but they do not inherently require manual endpoint management. A ClusterIP Service with selectors (D) is the standard case where endpoints are automatically created and updated. NodePort and LoadBalancer Services also typically use selectors and therefore inherit automatic endpoint management; the difference is in how traffic enters the cluster, not how backends are discovered.

Operationally, when using Services without selectors, you must ensure endpoint IPs remain correct, health is accounted for (often via external tooling), and you update endpoints when backends change. The key concept is: no selector → Kubernetes can't auto-populate endpoints → you must provide them.

Question: 63

Which of these commands is used to retrieve the documentation and field definitions for a Kubernetes resource?

- A. `kubectl explain`
- B. `kubectl api-resources`
- C. `kubectl get --help`
- D. `kubectl show`

Answer: A

Explanation:

`kubectl explain` is the command that shows documentation and field definitions for Kubernetes resource schemas, so A is correct. Kubernetes resources have a structured schema: top-level fields like `apiVersion`, `kind`, and `metadata`, and resource-specific structures like `spec` and `status`. `kubectl explain` lets you explore these structures directly from your cluster's API discovery information, including field types, descriptions, and nested fields.

For example, `kubectl explain deployment` describes the Deployment resource, and `kubectl explain deployment.spec` dives

into the spec structure. You can continue deeper, such as `kubectl explain deployment.spec.template.spec.containers` to discover container fields. This is especially useful when writing or troubleshooting manifests, because it reduces guesswork and prevents invalid YAML fields that would be rejected by the API server. It also helps when APIs evolve: you can confirm which fields exist in your cluster's current version and what they mean.

The other commands do different things. `kubectl api-resources` lists resource types and their shortnames, whether they are namespaced, and supported verbs—useful discovery, but not detailed field definitions. `kubectl get --help` shows CLI usage help for `kubectl get`, not the Kubernetes object schema. `kubectl show` is not a standard `kubectl` subcommand.

From a Kubernetes “declarative configuration” perspective, correct manifests are critical: controllers reconcile desired state from spec, and subtle field mistakes can change runtime behavior. `kubectl explain` is a built-in way to learn the schema and write manifests that align with the Kubernetes API's expectations. That's why it's commonly recommended in Kubernetes documentation and troubleshooting workflows.

Question: 64

Which of the following is a lightweight tool that manages traffic flows between services, enforces access policies, and aggregates telemetry data, all without requiring changes to application code?

- A. NetworkPolicy
- B. Linkerd
- C. kube-proxy
- D. Nginx

Answer: B

Explanation:

Linkerd is a lightweight service mesh that manages service-to-service traffic, security policies, and telemetry without requiring application code changes—so B is correct. A service mesh introduces a dedicated layer for east-west traffic (internal service calls) and typically provides features like mutual TLS (mTLS), retries/timeouts, traffic shaping, and consistent metrics/tracing signals. Linkerd is known for being simpler and resource-efficient relative to some alternatives, which aligns with the “lightweight tool” phrasing.

Why this matches the description: In a service mesh, workload traffic is intercepted by a proxy layer (often as a sidecar or node-level/ambient proxy) and managed centrally by mesh control components. This allows security and traffic policy to be applied uniformly without modifying each microservice. Telemetry is also generated consistently because the proxies observe traffic directly and emit metrics and traces about request rates, latency, and errors.

The other choices don't fit. NetworkPolicy is a Kubernetes resource that controls allowed network flows (L3/L4) but does not provide L7 traffic management, retries, identity-based mTLS, or automatic telemetry aggregation. kube-proxy implements Service networking rules (ClusterIP/NodePort forwarding) but does not enforce access policies at the service identity level and is not a telemetry system. Nginx can be used as an ingress controller or reverse proxy, but it is not inherently a full service mesh spanning all service-to-service communication and policy/telemetry across the mesh by default.

In cloud native architecture, service meshes help address cross-cutting concerns—security, observability, and traffic management—without embedding that logic into every application. The question's combination of "traffic flows," "access policies," and "aggregates telemetry" maps directly to a mesh, and the lightweight mesh option provided is Linkerd.

Question: 65

Which Kubernetes resource uses immutable: true boolean field?

- A. Deployment
- B. Pod
- C. ConfigMap
- D. ReplicaSet

Answer: C

Explanation:

The immutable: true field is supported by ConfigMap (and also by Secrets, though Secret is not in the options), so C is correct. When a ConfigMap is marked immutable, its data can no longer be changed after creation. This is useful for protecting configuration from accidental modification and for improving cluster performance by reducing watch/update churn on frequently referenced configuration objects.

In Kubernetes, ConfigMaps store non-sensitive configuration as key-value pairs. They can be consumed by Pods as

environment variables, command-line arguments, or mounted files in volumes. Without immutability, ConfigMap updates can trigger complex runtime behaviors: for example, filemounted ConfigMap updates can eventually reflect in the volume (with some delay), but environment variables do not update automatically in running Pods. This can cause confusion and configuration drift between expected and actual behavior. Marking a ConfigMap immutable makes the configuration stable and encourages explicit rollout strategies (create a new ConfigMap with a new name and update the Pod template), which is generally more reliable for production delivery.

Why the other options are wrong: Deployments, Pods, and ReplicaSets do not use an immutable: true field as a standard top-level toggle in their API schema for the purpose described. These objects can be updated through the normal API mechanisms, and their updates are part of typical lifecycle operations (rolling updates, scaling, etc.). The immutability concept exists in Kubernetes, but the specific immutable boolean in this context is a recognized field for ConfigMap (and Secret) objects.

Operationally, immutable ConfigMaps help enforce safer practices: instead of editing live configuration in place, teams adopt versioned configuration artifacts and controlled rollouts via Deployments. This fits cloud-native principles of repeatability and reducing accidental production changes.

Question: 66

Which statement about the Kubernetes network model is correct?

- A. Pods can only communicate with Pods exposed via a Service.
- B. Pods can communicate with all Pods without NAT.
- C. The Pod IP is only visible inside a Pod.
- D. The Service IP is used for the communication between Services.

Answer: B

Explanation:

Kubernetes' networking model assumes that every Pod has its own IP address and that Pods can communicate with other Pods across nodes without requiring network address translation (NAT). That makes B correct. This is one of Kubernetes' core design assumptions and is typically implemented via CNI plugins that provide flat, routable Pod networking (or equivalent behavior using encapsulation/routing).

This model matters because scheduling is dynamic. The scheduler can place Pods anywhere in the cluster, and applications should not need to know whether a peer is on the same node or a different node. With the Kubernetes

network model, Pod-to-Pod communication works uniformly: a Pod can reach any other Pod IP directly, and nodes can reach Pods as well. Services and DNS add stable naming and load balancing, but direct Pod connectivity is part of the baseline model.

Option A is incorrect because Pods can communicate directly using Pod IPs even without Services (subject to NetworkPolicies and routing). Services are abstractions for stable access and load balancing; they are not the only way Pods can communicate. Option C is incorrect because Pod IPs are not limited to visibility “inside a Pod”; they are routable within the cluster network. Option D is misleading: Services are often used by Pods (clients) to reach a set of Pods (backends). “Service IP used for communication between Services” is not the fundamental model; Services are virtual IPs for reaching workloads, and “Service-to-Service communication” usually means one workload calling another via the target Service name.

A useful way to remember the official model: (1) all Pods can communicate with all other Pods (no NAT), (2) all nodes can communicate with all Pods (no NAT), (3) Pod IPs are unique cluster-wide. This enables consistent microservice connectivity and supports higher-level traffic management layers like Ingress and service meshes.

Question: 67

What is the resource type used to package sets of containers for scheduling in a cluster?

- A. Pod
- B. ContainerSet
- C. ReplicaSet
- D. Deployment

Answer: A

Explanation:

The Kubernetes resource used to package one or more containers into a schedulable unit is the Pod, so A is correct. Kubernetes schedules Pods onto nodes; it does not schedule individual containers. A Pod represents a single “instance” of an application component and includes one or more containers that share key runtime properties, including the same network namespace (same IP and port space) and the ability to share volumes.

Pods enable common patterns beyond “one container per Pod.” For example, a Pod may include a main application container plus a sidecar container for logging, proxying, or configuration reload. Because these containers share localhost networking and volume mounts, they can coordinate efficiently without requiring external service calls. Kubernetes manages the Pod lifecycle as a unit: the containers in a Pod are started according to container lifecycle rules and are co-located on the same node.

Option B (ContainerSet) is not a standard Kubernetes workload resource. Option C (ReplicaSet) manages a set of Pod replicas, ensuring a desired count is running, but it is not the packaging unit itself. Option D (Deployment) is a higher-level controller that manages ReplicaSets and provides rollout/rollback behavior, again operating on Pods rather than being the container-packaging unit.

From the scheduling perspective, the PodSpec defines container images, commands, resources, volumes, security context, and placement constraints. The scheduler evaluates these constraints and assigns the Pod to a node. This “Pod as the atomic scheduling unit” is fundamental to Kubernetes architecture and explains why Kubernetes-native concepts (Services, selectors, readiness, autoscaling) all revolve around Pods.

Question: 68

Can a Kubernetes Service expose multiple ports?

- A. No, you can only expose one port per each Service.
- B. Yes, but you must specify an unambiguous name for each port.
- C. Yes, the only requirement is to use different port numbers.
- D. No, because the only port you can expose is port number 443.

Answer: B

Explanation:

Yes, a Kubernetes Service can expose multiple ports, and when it does, each port should have a unique, unambiguous name, making B correct. In the Service spec, the ports field is an array, allowing you to define multiple port mappings (e.g., 80 for HTTP and 443 for HTTPS, or grpc and metrics). Each entry can include port (Service port), targetPort (backend Pod port), and protocol.

The naming requirement becomes important because Kubernetes needs to disambiguate ports, especially when other resources refer to them. For example, an Ingress backend or some proxies/controllers can reference Service ports by name. Also, when multiple ports exist, a name helps humans and automation reliably select the correct port. Kubernetes documentation and common practice recommend naming ports whenever there is more than one, and in several scenarios it's effectively required to avoid ambiguity.

Option A is incorrect because multi-port Services are common and fully supported. Option C is insufficient: while different port numbers are necessary, naming is the correct distinguishing rule emphasized by Kubernetes patterns and required by some integrations. Option D is incorrect and nonsensical—Services can expose many ports and are not restricted to 443.

Operationally, exposing multiple ports through one Service is useful when a single backend workload provides multiple interfaces (e.g., application traffic and a metrics endpoint). You can keep stable discovery under one DNS name while still differentiating ports. The backend Pods must still listen on the target ports, and selectors determine which Pods are endpoints. The key correctness point for this question is: multi-port Services are allowed, and each port should be uniquely named to avoid confusion and integration issues.

Question: 69

Which persona is normally responsible for defining, testing, and running an incident management process?

- A. Site Reliability Engineers
- B. Project Managers
- C. Application Developers
- D. Quality Engineers

Answer: A

Explanation:

The role most commonly responsible for defining, testing, and running an incident management process is Site Reliability Engineers (SREs), so A is correct. SRE is an operational engineering discipline focused on ensuring reliability, availability, and performance of services in production. Incident management is a core part of that mission: when outages or severe degradations occur, someone must coordinate response, restore service quickly, and then drive follow-up improvements

to prevent recurrence.

In cloud native environments (including Kubernetes), incident response involves both technical and process elements. On the technical side, SREs ensure observability is in place—metrics, logs, traces, dashboards, and actionable alerts—so incidents can be detected and diagnosed quickly. They also validate operational readiness: runbooks, escalation paths, on-call rotations, and post-incident review practices. On the process side, SREs often establish severity classifications, response roles (incident commander, communications lead, subject matter experts), and “game day” exercises or simulated incidents to test preparedness.

Project managers may help coordinate schedules and communication for projects, but they are not typically the owners of operational incident response mechanics. Application developers are crucial participants during incidents, especially for debugging application-level failures, but they are not usually the primary maintainers of the incident management framework. Quality engineers focus on testing and quality assurance, and while they contribute to preventing defects, they are not usually the owners of real-time incident operations.

In Kubernetes specifically, incidents often span multiple layers: workload behavior, cluster resources, networking, storage, and platform dependencies. SREs are positioned to manage the cross-cutting operational view and to continuously improve reliability through error budgets, SLOs/SLIs, and iterative hardening. That’s why the correct persona is Site Reliability Engineers.

Question: 70

What is the default deployment strategy in Kubernetes?

- A. Rolling update
- B. Blue/Green deployment
- C. Canary deployment
- D. Recreate deployment

Answer: A

Explanation:

For Kubernetes Deployments, the default update strategy is RollingUpdate, which corresponds to “Rolling update” in option A. Rolling updates replace old Pods with new Pods gradually, aiming to maintain availability during the rollout.

Kubernetes does this by creating a new ReplicaSet for the updated Pod template and then scaling the new ReplicaSet up

while scaling the old one down.

The pace and safety of a rolling update are controlled by parameters like `maxUnavailable` and `maxSurge`. `maxUnavailable` limits how many replicas can be unavailable during the update, protecting availability. `maxSurge` controls how many extra replicas can be created temporarily above the desired count, helping speed up rollouts while maintaining capacity. If readiness probes fail, Kubernetes will pause progression because new Pods aren't becoming Ready, helping prevent a bad version from fully replacing a good one.

Options B (Blue/Green) and C (Canary) are popular progressive delivery patterns, but they are not the default built-in Deployment strategy. They are typically implemented using additional tooling (service mesh routing, traffic splitting controllers, or specialized rollout controllers) or by operating multiple Deployments/Services. Option D (Recreate) is a valid strategy but not the default; it terminates all old Pods before creating new ones, causing downtime unless you have external buffering or multi-tier redundancy.

From an application delivery perspective, RollingUpdate aligns with Kubernetes' declarative model: you update the desired Pod template and let the controller converge safely. `kubectl rollout status` is commonly used to monitor progress.

Rollbacks are also supported because the Deployment tracks history. Therefore, the verified correct answer is

A: Rolling update.

Question: 71

Which command provides information about the field replicas within the spec resource of a deployment object?

- A. `kubectl get deployment.spec.replicas`
- B. `kubectl explain deployment.spec.replicas`
- C. `kubectl describe deployment.spec.replicas`
- D. `kubectl explain deployment --spec.replicas`

Answer: B

Explanation:

The correct command to get field-level schema information about `spec.replicas` in a Deployment is `kubectl explain deployment.spec.replicas`, so B is correct. `kubectl explain` is designed to retrieve documentation for resource fields directly from Kubernetes API discovery and OpenAPI schemas. When you use `kubectl explain deployment.spec.replicas`,

kubectl shows what the field means, its type, and any relevant notes—exactly what “provides information about the field” implies.

This differs from `kubectl get` and `kubectl describe`. `kubectl get` is for retrieving actual objects or listing resources; it does not accept dot-paths like `deployment.spec.replicas` as a normal resource argument. You can use JSONPath/custom-columns with `kubectl get deployment <name> -o jsonpath='{.spec.replicas}'` to extract the live value, but that’s not what option A shows, and it’s not “documentation.” `kubectl describe` provides a human-friendly summary of a live resource instance (events, status, spec), but it doesn’t target schema field definitions via dot-path in the way shown in option C.

Option D is not valid syntax: `kubectl explain deployment --spec.replicas` is not how `kubectl explain` accepts nested field references. The correct pattern is positional dot notation: `kubectl explain <resource>.<field>.<subfield>`

Understanding `spec.replicas` matters operationally: it defines the desired number of Pod replicas for a Deployment. The Deployment controller ensures that the corresponding ReplicaSet maintains that count, supporting self-healing if Pods fail. While autoscalers can adjust replicas automatically, the field remains the primary declarative knob. The question is specifically about finding information (schema docs) for that field, which is why `kubectl explain deployment.spec.replicas` is the verified correct answer.

Question: 72

Which of the following is a responsibility of the governance board of an open source project?

- A. Decide about the marketing strategy of the project.
- B. Review the pull requests in the main branch.
- C. Outline the project’s “terms of engagement”.
- D. Define the license to be used in the project.

Answer: C

Explanation:

A governance board in an open source project typically defines how the community operates—its decision-making rules, roles, conflict resolution, and contribution expectations—so C (“Outline the project’s terms of engagement”) is correct. In large cloud-native projects (Kubernetes being a prime example), clear governance is essential to coordinate many

contributors, companies, and stakeholders. Governance establishes the “rules of the road” that keep collaboration productive and fair.

“Terms of engagement” commonly includes: how maintainers are selected, how proposals are reviewed (e.g., enhancement processes), how meetings and SIGs operate, what constitutes consensus, how voting works when consensus fails, and what code-of-conduct expectations apply. It also defines escalation and dispute resolution paths so technical disagreements don’t become community-breaking conflicts. In other words, governance is about ensuring the project has durable, transparent processes that outlive any individual contributor and support vendor-neutral decision making.

Option B (reviewing pull requests) is usually the responsibility of maintainers and SIG owners, not a governance board. The governance body may define the structure that empowers maintainers, but it generally does not do day-to-day code review. Option A (marketing strategy) is often handled by foundations, steering committees, or separate outreach groups, not governance boards as their primary responsibility. Option D (defining the license) is usually decided early and may be influenced by a foundation or legal process; while governance can shape legal/policy direction, the core governance responsibility is broader community operating rules rather than selecting a license.

In cloud-native ecosystems, strong governance supports sustainability: it encourages contributions, protects neutrality, and provides predictable processes for evolution. Therefore, the best verified answer is C.

Question: 73

What is the role of a NetworkPolicy in Kubernetes?

- A. The ability to cryptic and obscure all traffic.
- B. The ability to classify the Pods as isolated and non isolated.
- C. The ability to prevent loopback or incoming host traffic.
- D. The ability to log network security events.

Answer: B

Explanation:

A Kubernetes NetworkPolicy defines which traffic is allowed to and from Pods by selecting Pods and specifying ingress/egress rules. A key conceptual effect is that it can make Pods “isolated” (default deny except what is allowed) versus “non-isolated” (default allow). This aligns best with option B, so B is correct.

By default, Kubernetes networking is permissive: Pods can typically talk to any other Pod. When you apply a NetworkPolicy that selects a set of Pods, those selected Pods become “isolated” for the direction(s) covered by the policy (ingress and/or egress). That means only traffic explicitly allowed by the policy is permitted; everything else is denied (again, for the selected Pods and direction). This classification concept—isolated vs non-isolated—is a common way the Kubernetes documentation explains NetworkPolicy behavior.

Option A is incorrect: NetworkPolicy does not encrypt (“cryptic and obscure”) traffic. Encryption is typically handled by mTLS via a service mesh or application-layer TLS. Option C is not the primary role; loopback and host traffic handling depend on the network plugin and node configuration, and NetworkPolicy is not a “prevent loopback” mechanism. Option D is incorrect because NetworkPolicy is not a logging system; while some CNIs can produce logs about policy decisions, logging is not NetworkPolicy’s role in the API.

One critical Kubernetes detail: NetworkPolicy enforcement is performed by the CNI/network plugin. If your CNI doesn’t implement NetworkPolicy, creating these objects won’t change runtime traffic. In CNIs that do support it, NetworkPolicy becomes a foundational security primitive for segmentation and least privilege: restricting database access to app Pods only, isolating namespaces, and reducing lateral movement risk.

So, in the language of the provided answers, NetworkPolicy’s role is best captured as the ability to classify Pods into isolated/non-isolated by applying traffic-allow rules—option B.

Question: 74

What are the most important resources to guarantee the performance of an etcd cluster?

- A. CPU and disk capacity.
- B. Network throughput and disk I/O.
- C. CPU and RAM memory.
- D. Network throughput and CPU.

Answer: B

Explanation:

etcd is the strongly consistent key-value store backing Kubernetes cluster state. Its performance directly affects the entire control plane because most API operations require reads/writes to etcd. The most critical resources for etcd performance are disk I/O (especially latency) and network throughput/latency between etcd members and API servers—so B is correct.

etcd is write-ahead-log (WAL) based and relies heavily on stable, low-latency storage. Slow disks increase commit latency, which slows down object updates, watches, and controller loops. In busy clusters, poor disk performance can cause request backlogs and timeouts, showing up as slow kubectl operations and delayed controller reconciliation. That's why production guidance commonly emphasizes fast SSD-backed storage and careful monitoring of fsync latency.

Network performance matters because etcd uses the Raft consensus protocol. Writes must be replicated to a quorum of members, and leader-follower communication is continuous. High network latency or low throughput can slow replication and increase the time to commit writes. Unreliable networking can also cause leader elections or cluster instability, further degrading performance and availability.

CPU and memory are still relevant, but they are usually not the first bottleneck compared to disk and network. CPU affects request processing and encryption overhead if enabled, while memory affects caching and compaction behavior.

Disk "capacity" alone (size) is less relevant than disk I/O

characteristics (latency, IOPS), because etcd performance is sensitive to fsync and write latency.

In Kubernetes operations, ensuring etcd health includes: using dedicated fast disks, keeping network stable, enabling regular compaction/defragmentation strategies where appropriate, sizing correctly (typically odd-numbered members for quorum), and monitoring key metrics (commit latency, fsync duration, leader changes). Because etcd is the persistence layer of the API, disk I/O and network quality are the primary determinants of control-plane responsiveness—hence B.

Question: 75

How do you deploy a workload to Kubernetes without additional tools?

- A. Create a Bash script and run it on a worker node.
- B. Create a Helm Chart and install it with helm.
- C. Create a manifest and apply it with kubectl.
- D. Create a Python script and run it with kubectl.

Answer: C

Explanation:

The standard way to deploy workloads to Kubernetes using only built-in tooling is to create Kubernetes manifests (YAML/JSON definitions of API objects) and apply them with kubectl, so C is correct. Kubernetes is a declarative system: you describe the desired state of resources (e.g., a Deployment, Service, ConfigMap, Ingress) in a manifest file, then submit that desired state to the API server. Controllers reconcile the actual cluster state to match what you declared.

A manifest typically includes mandatory fields like apiVersion, kind, and metadata, and then a spec describing desired behavior. For example, a Deployment manifest declares replicas and the Pod template (containers, images, ports, probes, resources). Applying the manifest with kubectl apply -f <file> creates or updates the resources. kubectl apply is also designed to work well with iterative changes: you update the file, re-apply, and Kubernetes performs a controlled rollout based on controller logic.

Option B (Helm) is indeed a popular deployment tool, but Helm is explicitly an “additional tool” beyond kubectl and the Kubernetes API. The question asks “without additional tools,” so Helm is excluded by definition. Option A (running Bash scripts on worker nodes) bypasses Kubernetes’ desired-state control and is not how Kubernetes workload deployment is intended; it also breaks portability and operational safety. Option D is not a standard Kubernetes deployment mechanism; kubectl does not “run Python scripts” to deploy workloads (though scripts can automate kubectl, that’s still not the primary mechanism).

From a cloud native delivery standpoint, manifests support GitOps, reviewable changes, and repeatable deployments across environments. The Kubernetes-native approach is: declare resources in manifests and apply them to the cluster.

Therefore, C is the verified correct answer.

Question: 76

How do you perform a command in a running container of a Pod?

- A. kubectl exec <pod> -- <command>
- B. docker exec <pod> <command>
- C. kubectl run <pod> -- <command>
- D. kubectl attach <pod> -i <command>

Answer: A

Explanation:

In Kubernetes, the standard way to execute a command inside a running container is kubectl exec, which is why A is

correct. `kubectl exec` calls the Kubernetes API (API server), which then coordinates with the kubelet on the target node to run the requested command inside the container using the container runtime's exec mechanism. The `--` separator is important: it tells `kubectl` that everything after `--` is the command to run in the container rather than flags for `kubectl` itself.

This is fundamentally different from `docker exec`. In Kubernetes, you don't normally target containers through Docker/CRI tools directly because Kubernetes abstracts the runtime behind CRI. Also, "Docker" might not even be installed on nodes in modern clusters (`containerd`/`CRI-O` are common).

So option B is not the Kubernetes-native approach and often won't work.

`kubectl run` (option C) is for creating a new Pod (or generating workload resources), not for executing a command in an existing container. `kubectl attach` (option D) attaches your terminal to a running container's process streams (`stdin/stdout/stderr`), which is useful for interactive sessions, but it does not execute an arbitrary new command like `exec` does.

In real usage, you often specify the container when a Pod has multiple containers: `kubectl exec -it <pod> -c <container> -- /bin/sh`. This is common for debugging, verifying config files mounted from `ConfigMaps`/`Secrets`, testing DNS resolution, or checking network connectivity from within the Pod network namespace. Because `exec` uses the API and kubelet, it respects Kubernetes access control (RBAC) and audit logging—another reason it's the correct operational method.

Question: 77

How to create a headless Service?

- A. By specifying `.spec.clusterIP: headless`
- B. By specifying `.spec.clusterIP: None`
- C. By specifying `.spec.clusterIP: 0.0.0.0`
- D. By specifying `.spec.clusterIP: localhost`

Answer: B

Explanation:

A headless Service is created by setting `spec.clusterIP: None`, so B is correct. Normally, a Service gets a ClusterIP, and

kube-proxy (or an alternative dataplane) implements virtual-IP-based load balancing to route traffic from that ClusterIP to the backend Pods. A headless Service intentionally disables that virtual IP allocation. Instead of giving you a single stable VIP, Kubernetes publishes DNS records that resolve directly to the endpoints (the Pod IPs) behind the Service.

This is especially important for workloads that need direct endpoint discovery or stable per-Pod

identities, such as StatefulSets. With a headless Service, clients can discover all Pod IPs (or individual Pod DNS names in StatefulSet patterns) and implement their own selection, quorum, or leader/follower logic. Kubernetes DNS (CoreDNS) responds differently for headless Services: rather than returning a single ClusterIP, it returns multiple A/AAAA records (one per endpoint) or SRV records for named ports, enabling richer service discovery behavior.

The other options are invalid. "headless" is not a magic value for clusterIP; the API expects either an actual IP address assigned by the cluster or the special literal None. 0.0.0.0 and localhost are not valid ways to request headless semantics.

Kubernetes uses None specifically to signal "do not allocate a ClusterIP."

Operationally, headless Services are used to: (1) expose each backend instance individually, (2) support stateful clustering and stable DNS names, and (3) avoid load balancing when the application or client library must choose endpoints itself.

The key is that the Service still provides a stable DNS name, but the resolution yields endpoints, not a VIP.

Question: 78

How does dynamic storage provisioning work?

- A. A user requests dynamically provisioned storage by including an existing StorageClass in their PersistentVolumeClaim.
- B. An administrator creates a StorageClass and includes it in their Pod YAML definition file without creating a PersistentVolumeClaim.
- C. A Pod requests dynamically provisioned storage by including a StorageClass and the Pod name in their PersistentVolumeClaim.
- D. An administrator creates a PersistentVolume and includes the name of the PersistentVolume in their Pod YAML definition file.

Answer: A

Explanation:

Dynamic provisioning is the Kubernetes mechanism where storage is created on-demand when a user creates a PersistentVolumeClaim (PVC) that references a StorageClass, so A is correct. In this model, the user does not need to pre-create a PersistentVolume (PV). Instead, the StorageClass points to a provisioner (typically a CSI driver) that knows how to create a volume in the underlying storage system (cloud disk, SAN, NAS, etc.). When the PVC is created with storageClassName: <class>, Kubernetes triggers the provisioner to create a new volume and then binds the resulting PV to that PVC.

This is why option B is incorrect: you do not put a StorageClass "in the Pod YAML" to request provisioning. Pods reference PVCs, not StorageClasses directly. Option C is incorrect because the PVC does not need the Pod name; binding is done via the PVC itself. Option D describes static provisioning: an admin pre-creates PVs and users claim them by creating PVCs that match the PV (capacity, access modes, selectors). Static provisioning can work, but it is not dynamic provisioning.

Under the hood, the StorageClass can define parameters like volume type, replication, encryption, and binding behavior (e.g., volumeBindingMode: WaitForFirstConsumer to delay provisioning until the Pod is scheduled, ensuring the volume is created in the correct zone). Reclaim policies (Delete/Retain) define what happens to the underlying volume after the PVC is deleted.

In cloud-native operations, dynamic provisioning is preferred because it improves developer selfservice, reduces manual admin work, and makes scaling stateful workloads easier and faster. The essence is: PVC + StorageClass → automatic PV creation and binding.

Question: 79

Which of the following are tasks performed by a container orchestration tool?

- A. Schedule, scale, and manage the health of containers.
- B. Create images, scale, and manage the health of containers.
- C. Debug applications, and manage the health of containers.
- D. Store images, scale, and manage the health of containers.

Answer: A

Explanation:

A container orchestration tool (like Kubernetes) is responsible for scheduling, scaling, and health management of workloads, making A correct. Orchestration sits above individual containers and focuses on running applications reliably across a fleet of machines. Scheduling means deciding which node should run a workload based on resource requests, constraints, affinities, taints/tolerations, and current cluster state. Scaling means changing the number of running instances (replicas) to meet demand (manually or automatically through autoscalers). Health management includes monitoring whether containers and Pods are alive and ready, replacing failed instances, and maintaining the declared desired state.

Options B and D include “create images” and “store images,” which are not orchestration responsibilities. Image creation is a CI/build responsibility (Docker/BuildKit/build systems), and image storage is a container registry responsibility (Harbor, ECR, GCR, Docker Hub, etc.). Kubernetes consumes images from registries but does not build or store them.

Option C includes “debug applications,” which is not a core orchestration function. While Kubernetes provides tools that help debugging (logs, exec, events), debugging is a human/operator activity rather than the orchestrator’s fundamental responsibility.

In Kubernetes specifically, these orchestration tasks are implemented through controllers and control loops:

Deployments/ReplicaSets manage replica counts and rollouts, kube-scheduler assigns Pods to nodes, kubelet ensures containers run, and probes plus controller logic replace unhealthy replicas. This is exactly what makes Kubernetes valuable at scale: instead of manually starting/stopping containers on individual hosts, you declare your intent and let the orchestration system continually reconcile reality to match. That combination—placement + elasticity + self-healing—is the core of container orchestration, matching option A precisely.

Question: 80

Which of the following is a definition of Hybrid Cloud?

- A. A combination of services running in public and private data centers, only including data centers from the same cloud provider.
- B. A cloud native architecture that uses services running in public clouds, excluding data centers in different availability zones.
- C. A cloud native architecture that uses services running in different public and private clouds, including on-premises data centers.

D. A combination of services running in public and private data centers, excluding serverless functions.

Answer: C

Explanation:

A hybrid cloud architecture combines public cloud and private/on-premises environments, often spanning multiple infrastructure domains while maintaining some level of portability, connectivity, and unified operations. Option C captures the commonly accepted definition: services run across public and private clouds, including on-premises data centers, so C is correct.

Hybrid cloud is not limited to a single cloud provider (which is why A is too restrictive). Many organizations adopt hybrid cloud to meet regulatory requirements, data residency constraints, latency needs, or to preserve existing investments while still using public cloud elasticity. In Kubernetes terms, hybrid strategies often include running clusters both on-prem and in one or more public clouds, then standardizing deployment through Kubernetes APIs, GitOps, and consistent security/observability practices.

Option B is incorrect because excluding data centers in different availability zones is not a defining property; in fact, hybrid deployments commonly use multiple zones/regions for resilience. Option D is a distraction: serverless inclusion or exclusion does not define hybrid cloud. Hybrid is about the combination of infrastructure environments, not a specific compute model.

A practical cloud-native view is that hybrid architectures introduce challenges around identity, networking, policy enforcement, and consistent observability across environments. Kubernetes helps because it provides a consistent control plane API and workload model regardless of where it runs. Tools like service meshes, federated identity, and unified monitoring can further reduce fragmentation.

So, the most accurate definition in the given choices is C: hybrid cloud combines public and private clouds, including on-premises infrastructure, to run services in a coordinated architecture.

Question: 81

What is a Kubernetes Service Endpoint?

A. It is the API endpoint of our Kubernetes cluster.

B. It is a name of special Pod in kube-system namespace.

- C. It is an IP address that we can access from the Internet.
- D. It is an object that gets IP addresses of individual Pods assigned to it.

Answer: D

Explanation:

A Kubernetes Service routes traffic to a dynamic set of backends (usually Pods). The set of backend IPs and ports is represented by endpoint-tracking resources. Historically this was the Endpoints object; today Kubernetes commonly uses EndpointSlice for scalability, but the concept remains the same: endpoints represent the concrete network destinations behind a Service. That's why D is correct: a Service endpoint is an object that contains the IP addresses (and ports) of the individual Pods (or other backends) associated with that Service.

When a Service has a selector, Kubernetes automatically maintains endpoints by watching which Pods match the selector and are Ready, then publishing those Pod IPs into Endpoints/EndpointSlices. Consumers don't usually use endpoints directly; instead they call the Service DNS name, and kube-proxy (or an alternate dataplane) forwards traffic to one of the endpoints. Still, endpoints are critical because they are what make Service routing accurate and up to date during scaling events, rolling updates, and failures.

Option A confuses this with the Kubernetes API server endpoint (the cluster API URL). Option B is incorrect; there's no special "Service Endpoint Pod." Option C describes an external/public IP concept, which may exist for LoadBalancer Services, but "Service endpoint" in Kubernetes vocabulary is about the backend destinations, not the public endpoint.

Operationally, endpoints are useful for debugging: if a Service isn't routing traffic, checking Endpoints/EndpointSlices shows whether the Service actually has backends and whether readiness is excluding Pods. This ties directly into Kubernetes service discovery and load balancing: the Service is the stable front door; endpoints are the actual backends.

Question: 82

Why is Cloud-Native Architecture important?

- A. Cloud Native Architecture revolves around containers, microservices and pipelines.
- B. Cloud Native Architecture removes constraints to rapid innovation.
- C. Cloud Native Architecture is modern for application deployment and pipelines.

D. Cloud Native Architecture is a bleeding edge technology and service.

Answer: B

Explanation:

Cloud-native architecture is important because it enables organizations to build and run software in a way that supports rapid innovation while maintaining reliability, scalability, and efficient operations. Option B best captures this: cloud native removes constraints to rapid innovation, so B is correct.

In traditional environments, innovation is slowed by heavyweight release processes, tightly coupled systems, manual operations, and limited elasticity. Cloud-native approaches—containers, declarative APIs, automation, and microservices-friendly patterns—reduce those constraints. Kubernetes exemplifies this by offering a consistent deployment model, self-healing, automated rollouts, scaling primitives, and a large ecosystem of delivery and observability tools. This makes it easier to ship changes more frequently and safely: teams can iterate quickly, roll back confidently, and standardize operations across environments.

Option A is partly descriptive (containers/microservices/pipelines are common in cloud native), but it doesn't explain why it matters; it lists ingredients rather than the benefit. Option C is vague ("modern") and again doesn't capture the core value proposition. Option D is incorrect because cloud native is not primarily about being "bleeding edge"—it's about proven practices that improve time-to-market and operational stability.

A good way to interpret "removes constraints" is: cloud native shifts the bottleneck away from infrastructure friction. With automation (IaC/GitOps), standardized runtime packaging (containers), and platform capabilities (Kubernetes controllers), teams spend less time on repetitive manual work

and more time delivering features. Combined with observability and policy automation, this results in faster delivery with better reliability—exactly the reason cloud-native architecture is emphasized across the Kubernetes ecosystem.

Question: 83

Which Kubernetes component is the smallest deployable unit of computing?

A. StatefulSet

B. Deployment

C. Pod

D. Container

Answer: C

Explanation:

In Kubernetes, the Pod is the smallest deployable and schedulable unit, making C correct. Kubernetes does not schedule individual containers directly; instead, it schedules Pods, each of which encapsulates one or more containers that must run together on the same node. This design supports both single-container Pods (the most common) and multi-container Pods (for sidecars, adapters, and co-located helper processes).

Pods provide shared context: containers in a Pod share the same network namespace (one IP address and port space) and can share storage volumes. This enables tight coupling where needed—for example, a service mesh proxy sidecar and the application container communicate via localhost, or a log-forwarding sidecar reads logs from a shared volume. Kubernetes manages lifecycle at the Pod level: kubelet ensures the containers defined in the PodSpec are running and uses probes to determine readiness and liveness.

StatefulSet and Deployment are controllers that manage sets of Pods. A Deployment manages ReplicaSets for stateless workloads and provides rollout/rollback features; a StatefulSet provides stable identities, ordered operations, and stable storage for stateful replicas. These are higher-level constructs, not the smallest units.

Option D (“Container”) is smaller in an abstract sense, but it is not the smallest Kubernetes deployable unit because Kubernetes APIs and scheduling work at the Pod boundary. You don’t “kubectl apply” a container; you apply a Pod template within a Pod object (often via controllers).

Understanding Pods as the atomic unit is crucial: Services select Pods, autoscalers scale Pods (replica counts), and scheduling decisions are made per Pod. That’s why Kubernetes documentation consistently refers to Pods as the fundamental building block for running workloads.

Question: 84

What kubectl command is used to retrieve the resource consumption (CPU and memory) for nodes or Pods?

- A. kubectl cluster-info
- B. kubectl version
- C. kubectl top
- D. kubectl api-resources

Answer: C

Explanation:

To retrieve CPU and memory consumption for nodes or Pods, you use `kubectl top`, so C is correct. `kubectl top nodes` shows per-node resource usage, and `kubectl top pods` shows per-Pod (and optionally per-container) usage. This data comes from the Kubernetes resource metrics pipeline, most commonly `metrics-server`, which scrapes `kubelet/cAdvisor` stats and exposes them via the `metrics.k8s.io` API.

It's important to recognize that `kubectl top` provides current resource usage snapshots, not long-term historical trending. For long-term metrics and alerting, clusters typically use Prometheus and related

tooling. But for quick operational checks—"Is this Pod CPU-bound?" "Are nodes near memory saturation?"—`kubectl top` is the built-in day-to-day tool.

Option A (`kubectl cluster-info`) shows general cluster endpoints and info about control plane services, not resource usage. Option B (`kubectl version`) prints client/server version info. Option D (`kubectl api-resources`) lists resource types available in the cluster. None of those report CPU/memory usage.

In observability practice, `kubectl top` is often used during incidents to correlate symptoms with resource pressure. For example, if a node is high on memory, you might see Pods being OOMKilled or the kubelet evicting Pods under pressure. Similarly, sustained high CPU utilization might explain latency spikes or throttling if limits are set. Note that `kubectl top` requires `metrics-server` (or an equivalent provider) to be installed and functioning; otherwise it may return errors like "metrics not available."

So, the correct command for retrieving node/Pod CPU and memory usage is `kubectl top`.

Question: 85

Which are the two primary modes for Service discovery within a Kubernetes cluster?

- A. Environment variables and DNS
- B. API calls and LDAP
- C. Labels and RADIUS
- D. Selectors and DHCP

Answer: A

Explanation:

Kubernetes supports two primary built-in modes of Service discovery for workloads: environment variables and DNS, making A correct.

Environment variables: When a Pod is created, kubelet can inject environment variables for Services

that exist in the same namespace at the time the Pod starts. These variables include the Service host and port (for example, MY_SERVICE_HOST and MY_SERVICE_PORT). This approach is simple but has limitations: values are captured at Pod creation time and don't automatically update if Services change, and it can become cluttered in namespaces with many Services.

DNS-based discovery: This is the most common and flexible method. Kubernetes cluster DNS (usually CoreDNS) provides names like service-name.namespace.svc.cluster.local. Clients resolve the name and connect to the Service, which then routes to backend Pods. DNS scales better, is dynamic with endpoint updates, supports headless Services for per-Pod discovery, and is the default pattern for microservice communication.

The other options are not Kubernetes service discovery modes. Labels and selectors are used internally to relate Services to Pods, but they are not what application code uses for discovery (apps typically don't query selectors; they call DNS names). LDAP and RADIUS are identity/authentication protocols, not service discovery. DHCP is for IP assignment on networks, not for Kubernetes Service discovery.

Operationally, DNS is central: many applications assume name-based connectivity. If CoreDNS is misconfigured or overloaded, service-to-service calls may fail even if Pods and Services are otherwise healthy. Environment-variable discovery can still work for some legacy apps, but modern cloudnative practice strongly prefers DNS (and sometimes service meshes on top of it). The key exam concept is: Kubernetes provides service discovery via env vars and DNS.

Question: 86

Which of the following capabilities are you allowed to add to a container using the Restricted policy?

- A. CHOWN
- B. SYS_CHROOT
- C. SETUID
- D. NET_BIND_SERVICE

Answer: D

Explanation:

Under the Kubernetes Pod Security Standards (PSS), the Restricted profile is the most locked-down baseline intended to reduce container privilege and host attack surface. In that profile, adding Linux capabilities is generally prohibited except for very limited cases. Among the listed capabilities, NET_BIND_SERVICE is the one commonly permitted in restricted-like policies, so D is correct.

NET_BIND_SERVICE allows a process to bind to “privileged” ports below 1024 (like 80/443) without running as root. This aligns with restricted security guidance: applications should run as non-root, but still sometimes need to listen on standard ports. Allowing NET_BIND_SERVICE enables that pattern without granting broad privileges.

The other capabilities listed are more sensitive and typically not allowed in a restricted profile: CHOWN can be used to change file ownership, SETUID relates to privilege changes and can be abused, and SYS_CHROOT is a broader system-level capability associated with filesystem root changes. In hardened Kubernetes environments, these are normally disallowed because they increase the risk of privilege escalation or container breakout paths, especially if combined with other misconfigurations.

A practical note: exact enforcement depends on the cluster’s admission configuration (e.g., the built-in Pod Security Admission controller) and any additional policy engines (OPA/Gatekeeper). But the security intent of “Restricted” is consistent: run as non-root, disallow privilege escalation, restrict capabilities, and lock down host access.

NET_BIND_SERVICE is a well-known exception used to support common application networking needs while staying non-root.

So, the verified correct choice for an allowed capability in Restricted among these options is D: NET_BIND_SERVICE.

Question: 87

What methods can you use to scale a Deployment?

- A. With kubectl edit deployment exclusively.
- B. With kubectl scale-up deployment exclusively.
- C. With kubectl scale deployment and kubectl edit deployment.
- D. With kubectl scale deployment exclusively.

Answer: C

Explanation:

A Deployment's replica count is controlled by `spec.replicas`. You can scale a Deployment by changing that field—either directly editing the object or using kubectl's scaling helper. Therefore C is correct: you can scale using `kubectl scale` and also via `kubectl edit`.

`kubectl scale deployment <name> --replicas=<n>` is the purpose-built command for scaling. It updates the Deployment's desired replica count and lets the Deployment controller and ReplicaSet reconcile the change by creating or deleting Pods. This is the cleanest and most explicit operational approach, and it's easy to automate in scripts and pipelines.

`kubectl edit deployment <name>` opens the live object in an editor, allowing you to modify fields such as `spec.replicas` manually. When you save and exit, kubectl submits the updated object back to the API server. This method is useful for quick interactive adjustments or when you're already making other spec edits, but it's less structured and more error-prone than `kubectl scale` for simple replica changes.

Option B is invalid because `kubectl scale-up deployment` is not a standard kubectl command. Option A is incorrect because `kubectl edit` is not the only method; scaling is commonly done with `kubectl scale`. Option D is also incorrect because while `kubectl scale` is a primary method, `kubectl edit` is also a valid method to change replicas.

In production, you often scale with autoscalers (HPA/VPA), but the question is asking about kubectl methods. The key Kubernetes concept is that scaling is achieved by updating desired state (`spec.replicas`), and controllers reconcile Pods to match.

Question: 88

What is a sidecar container?

- A. A Pod that runs next to another container within the same Pod.
- B. A container that runs next to another Pod within the same namespace.
- C. A container that runs next to another container within the same Pod.
- D. A Pod that runs next to another Pod within the same namespace.

Answer: C

Explanation:

A sidecar container is an additional container that runs alongside the main application container within the same Pod, sharing network and storage context. That matches option C, so C is correct. The sidecar pattern is used to add supporting capabilities to an application without modifying the application code. Because both containers are in the same Pod, the sidecar can communicate with the main container over localhost and share volumes for files, sockets, or logs.

Common sidecar examples include: log forwarders that tail application logs and ship them to a logging system, proxies (service mesh sidecars like Envoy) that handle mTLS and routing policy, config reloaders that watch ConfigMaps and signal the main process, and local caching agents. Sidecars are especially powerful in cloud-native systems because they standardize cross-cutting concerns— security, observability, traffic policy—across many workloads.

Options A and D incorrectly describe “a Pod running next to ...” which is not how sidecars work; sidecars are containers, not separate Pods. Running separate Pods “next to” each other in a namespace does not give the same shared network namespace and tightly coupled lifecycle. Option B is also incorrect for the same reason: a sidecar is not a separate Pod; it is a container in the same Pod.

Operationally, sidecars share the Pod lifecycle: they are scheduled together, scaled together, and generally terminated together. This is both a benefit (co-location guarantees) and a responsibility (resource requests/limits should include the sidecar’s needs, and failure modes should be understood). Kubernetes is increasingly formalizing sidecar behavior (e.g., sidecar containers with ordered startup semantics), but the core definition remains: a helper container in the same Pod.

Question: 89

Which is an industry-standard container runtime with an “emphasis” on simplicity, robustness, and portability?

- A. CRI-O
- B. LXD
- C. containerd
- D. kata-runtime

Answer: C

Explanation:

containerd is a widely adopted, industry-standard container runtime known for simplicity, robustness, and portability, so C is correct. containerd originated as a core component extracted from Docker and has become a common runtime across Kubernetes distributions and managed services. It implements container lifecycle management (image pull, unpack, container execution, snapshotting) and typically delegates low-level container execution to an OCI runtime like runc.

In Kubernetes, kubelet communicates with container runtimes through CRI. containerd provides a CRI plugin (or can be integrated via CRI implementations) that makes it a first-class choice for Kubernetes nodes. This aligns with the runtime landscape after dockershim removal: Kubernetes users commonly run containerd or CRI-O as the node runtime.

Option A (CRI-O) is also a CRI-focused runtime and is valid in Kubernetes contexts, but the phrasing “industry-standard ... emphasis on simplicity, robustness, and portability” is strongly associated with containerd’s positioning and broad cross-platform adoption beyond Kubernetes. Option B (LXD) is a system container manager (often associated with LXC) and not the standard Kubernetes runtime in mainstream CRI discussions. Option D (kata-runtime) is associated with Kata Containers, which focuses on stronger isolation by running containers inside lightweight VMs; that is a security-oriented sandbox approach rather than a simplicity/portability “industry standard” baseline runtime.

From a cloud-native operations point of view, containerd’s popularity comes from its stable API, strong ecosystem support, and alignment with OCI standards. It integrates cleanly with image registries, supports modern snapshotters, and is heavily used in production by many Kubernetes providers. Therefore, the best verified answer is C:

containerd.

Question: 90

What does vertical scaling an application deployment describe best?

- A. Adding/removing applications to meet demand.
- B. Adding/removing node instances to the cluster to meet demand.
- C. Adding/removing resources to applications to meet demand.
- D. Adding/removing application instances of the same application to meet demand.

Answer: C

Explanation:

Vertical scaling means changing the resources allocated to a single instance of an application (more or less CPU/memory), which is why C is correct. In Kubernetes terms, this corresponds to adjusting container resource requests and limits (for CPU and memory). Increasing resources can help a workload handle more load per Pod by giving it more compute or memory headroom; decreasing can reduce cost and improve cluster packing efficiency.

This differs from horizontal scaling, which changes the number of instances (replicas). Option D describes horizontal scaling: adding/removing replicas of the same workload, typically managed by a Deployment and often automated via the Horizontal Pod Autoscaler (HPA). Option B describes scaling the infrastructure layer (nodes) which is cluster/node autoscaling (Cluster Autoscaler in cloud environments). Option A is not a standard scaling definition.

In practice, vertical scaling in Kubernetes can be manual (edit the Deployment resource requests/limits) or automated using the Vertical Pod Autoscaler (VPA), which can recommend or apply new requests based on observed usage. A key nuance is that changing requests/limits often requires Pod restarts to take effect, so vertical scaling is less “instant” than HPA and can disrupt workloads if not planned. That’s why many production teams prefer horizontal scaling for traffic-driven workloads and use vertical scaling to right-size baseline resources or address memory- bound/cpu-bound behavior.

From a cloud-native architecture standpoint, understanding vertical vs horizontal scaling helps you

design for elasticity: use vertical scaling to tune per-instance capacity; use horizontal scaling for resilience and throughput; and combine with node autoscaling to ensure the cluster has sufficient capacity. The definition the question is testing is simple: vertical scaling = change resources per application instance, which is option C.

Question: 91

Which Prometheus metric represents a single value that can go up and down?

- A. Counter
- B. Gauge
- C. Summary
- D. Histogram

Answer: B

Explanation:

In Prometheus, a Gauge is the metric type used to represent a value that can increase and decrease over time, so B is correct. Gauges are suited for “current state” measurements such as current memory usage, number of active sessions, queue depth, temperature, or CPU usage—anything that can move up and down as the system changes.

This contrasts with a Counter (A), which is monotonically increasing (it only goes up, except when a process restarts and the counter resets to zero). Counters are ideal for totals like total HTTP requests served, total errors, or bytes sent, and you typically use `rate()`/`irate()` in PromQL to convert counters into per-second rates.

A Summary (C) and Histogram (D) are used for distributions, commonly request latency. Histograms record observations into buckets and can produce percentiles using functions like `histogram_quantile()`. Summaries compute quantiles on the client side and expose them directly, along with counts and sums. Neither of these is the simplest “single value that goes up and down” type.

In Kubernetes observability, Prometheus is often used to scrape metrics from cluster components

(API server, kubelet) and applications. Choosing the right metric type matters operationally: use gauges for instantaneous measurements, counters for event totals, and histograms/summaries for latency distributions. That’s why Prometheus documentation and best practices emphasize understanding metric semantics—because misusing types leads to incorrect alerts and dashboards.

So for a single numeric value that can go up and down, the correct metric type is Gauge, option B.

Question: 92

What is Serverless computing?

- A. A computing method of providing backend services on an as-used basis.
- B. A computing method of providing services for AI and ML operating systems.
- C. A computing method of providing services for quantum computing operating systems.
- D. A computing method of providing services for cloud computing operating systems.

Answer: A

Explanation:

Serverless computing is a cloud execution model where the provider manages infrastructure concerns and you consume compute as a service, typically billed based on actual usage (requests, execution time, memory), which matches A. In other words, you deploy code (functions) or sometimes containers, configure triggers (HTTP events, queues, schedules), and the platform automatically provisions capacity, scales it up/down, and handles much of availability and fault tolerance behind the scenes.

From a cloud-native architecture standpoint, “serverless” doesn’t mean there are no servers; it means developers don’t manage servers. The platform abstracts away node provisioning, OS patching, and much of runtime scaling logic. This aligns with the “as-used basis” phrasing: you pay for what you run rather than maintaining always-on capacity.

It’s also useful to distinguish serverless from Kubernetes. Kubernetes automates orchestration

(scheduling, self-healing, scaling), but operating Kubernetes still involves cluster-level capacity decisions, node pools, upgrades, networking baseline, and policy. With serverless, those responsibilities are pushed further toward the provider/platform. Kubernetes can enable serverless experiences (for example, event-driven autoscaling frameworks), but serverless as a model is about a higher level of abstraction than “orchestrate containers yourself.”

Options B, C, and D are incorrect because they describe specialized or vague “operating system” services rather than the commonly accepted definition. Serverless is not specifically about AI/ML OSs or quantum OSs; it’s a general compute delivery model that can host many kinds of workloads.

Therefore, the correct definition in this question is A: providing backend services on an as-used basis.

Question: 93

What is the purpose of the CRI?

- A. To provide runtime integration control when multiple runtimes are used.
- B. Support container replication and scaling on nodes.
- C. Provide an interface allowing Kubernetes to support pluggable container runtimes.
- D. Allow the definition of dynamic resource criteria across containers.

Answer: C

Explanation:

The Container Runtime Interface (CRI) exists so Kubernetes can support pluggable container runtimes behind a stable interface, which makes C correct. In Kubernetes, the kubelet is responsible for managing Pods on a node, but it does not implement container execution itself. Instead, it delegates container lifecycle operations (pull images, create pod sandbox, start/stop containers, fetch logs, exec/attach streaming) to a container runtime through a well-defined API. CRI is that API contract.

Because of CRI, Kubernetes can run with different container runtimes—commonly containerd or CRI-

O—without changing kubelet core logic. This improves portability and keeps Kubernetes modular: runtime innovation can happen independently while Kubernetes retains a consistent operational model. CRI is accessed via gRPC and defines the services and message formats kubelet uses to communicate with runtimes.

Option B is incorrect because replication and scaling are handled by controllers (Deployments/ReplicaSets) and schedulers, not by CRI. Option D is incorrect because resource criteria (requests/limits) are expressed in Pod specs and enforced via OS mechanisms (cgroups) and kubelet/runtime behavior, but CRI is not “for defining dynamic resource criteria.” Option A is vague and not the primary statement; while CRI enables runtime integration, its key purpose is explicitly to make runtimes pluggable and interoperable.

This design became even more important as Kubernetes moved away from Docker Engine integration (dockershim removal from kubelet). With CRI, Kubernetes focuses on orchestrating Pods, while runtimes focus on executing containers. That separation of responsibilities is a core container orchestration principle and is exactly what the question is testing.

So the verified answer is C.

Question: 94

Imagine there is a requirement to run a database backup every day. Which Kubernetes resource could be used to achieve that?

- A. kube-scheduler
- B. CronJob
- C. Task
- D. Job

Answer: B

Explanation:

To run a workload on a repeating schedule (like “every day”), Kubernetes provides CronJob, making B correct. A CronJob creates Jobs according to a cron-formatted schedule, and then each Job creates one or more Pods that run to completion. This is the Kubernetes-native replacement for traditional cron scheduling, but implemented as a declarative resource managed by controllers in the cluster.

For a daily database backup, you’d define a CronJob with a schedule (e.g., “0 2 * * *” for 2:00 AM daily), and specify the Pod template that performs the backup (invokes backup scripts/tools, writes output to durable storage, uploads to object storage, etc.). Kubernetes will then create a Job at each scheduled time. CronJobs also support operational controls like concurrencyPolicy (Allow/Forbid/Replace) to decide what happens if a previous backup is still running, startingDeadlineSeconds to handle missed schedules, and history limits to retain recent successful/failed Job records for debugging.

Option D (Job) is close but not sufficient for “every day.” A Job runs a workload until completion once; you would need an external scheduler to create a Job every day. Option A (kube-scheduler) is a control plane component responsible for placing Pods onto nodes and does not schedule recurring tasks. Option C (“Task”) is not a standard Kubernetes workload resource.

This question is fundamentally about mapping a recurring operational requirement (backup cadence) to Kubernetes primitives. The correct design is: CronJob triggers Job creation on a schedule; Job runs Pods to completion. Therefore, the correct answer is B.

Question: 95

In CNCF, who develops specifications for industry standards around container formats and runtimes?

- A. Open Container Initiative (OCI)
- B. Linux Foundation Certification Group (LFCG)
- C. Container Network Interface (CNI)
- D. Container Runtime Interface (CRI)

Answer: A

Explanation:

The organization responsible for defining widely adopted standards around container formats and runtime specifications is the Open Container Initiative (OCI), so A is correct. OCI defines the image specification (how container images are structured and stored) and the runtime specification (how to run a container), enabling interoperability across tooling and vendors. This is foundational to the cloud-native ecosystem because it allows different build tools, registries, runtimes, and orchestration platforms to work together reliably.

Within Kubernetes and CNCF-adjacent ecosystems, OCI standards are the reason an image built by one tool can be pushed to a registry and pulled/run by many different runtimes. For example, a Kubernetes node running containerd or CRI-O can run OCI-compliant images consistently. OCI standardization reduces fragmentation and vendor lock-in, which is a core motivation in open source cloud-native architecture.

The other options are not correct for this question. CNI (Container Network Interface) is a standard for configuring container networking, not container image formats and runtimes. CRI (Container Runtime Interface) is a Kubernetes-specific interface between kubelet and container runtimes—it enables pluggable runtimes for Kubernetes, but it is not the industry standard body for container format/runtime specifications. “LFCG” is not the recognized standards body here.

In short: OCI defines the “language” for container images and runtime behavior, which is why the same image can be executed across environments. Kubernetes relies on those standards indirectly through runtimes and tooling, but the specification work is owned by OCI. Therefore, the verified correct answer is A.

Question: 96

Which of the following options includes valid API versions?

- A. alpha1v1, beta3v3, v2
- B. alpha1, beta3, v2
- C. v1alpha1, v2beta3, v2
- D. v1alpha1, v2beta3, 2.0

Answer: C

Explanation:

Kubernetes API versions follow a consistent naming pattern that indicates stability level and versioning. The valid forms include stable versions like v1, and pre-release versions such as v1alpha1, v1beta1, etc. Option C contains valid-looking Kubernetes version strings—v1alpha1, v2beta3, v2—so C is correct.

In Kubernetes, the “v” prefix is part of the standard for API versions. A stable API uses v1, v2, etc. Prerelease APIs include a stability marker: alpha (earliest, most changeable) and beta (more stable but still may change). The numeric suffix (e.g., alpha1, beta3) indicates iteration within that stability stage.

Option A is invalid because strings like alpha1v1 and beta3v3 do not match Kubernetes conventions (the v comes first, and alpha/beta are qualifiers after the version: v1alpha1). Option B is invalid because alpha1 and beta3 are missing the leading version prefix; Kubernetes API versions are not just “alpha1.” Option D includes 2.0, which looks like semantic versioning but is not the Kubernetes API version format. Kubernetes uses v2, not 2.0, for API versions.

Understanding this matters because API versions signal compatibility guarantees. Stable APIs are supported for a defined deprecation window, while alpha/beta APIs may change in incompatible ways and can be removed more easily. When authoring manifests, selecting the correct apiVersion ensures the API server accepts your resource and that controllers interpret fields correctly.

Therefore, among the choices, C is the only option comprised of valid Kubernetes-style API version strings.

Question: 97

Which of the following will view the snapshot of previously terminated ruby container logs from Pod web-1?

- A. `kubectl logs -p -c ruby web-1`
- B. `kubectl logs -c ruby web-1`
- C. `kubectl logs -p ruby web-1`
- D. `kubectl logs -p -c web-1 ruby`

Answer: A

Explanation:

To view logs from the previously terminated instance of a container, you use `kubectl logs -p`. To select a specific container in a multi-container Pod, you use `-c <containerName>`. Combining both gives the correct command for "previous logs from the ruby container in Pod web-1," which is option A: `kubectl logs -p -c ruby web-1`.

The `-p` (or `--previous`) flag instructs `kubectl` to fetch logs for the prior container instance. This is most useful when the container has restarted due to a crash (`CrashLoopBackOff`) or was terminated and restarted. Without `-p`, `kubectl logs` shows logs for the currently running container instance (or the most recent if it's completed, depending on state).

Option B is close but wrong for the

question: it selects

the ruby container (`-c ruby`) but does not request the previous instance snapshot, so it returns current logs, not the prior-terminated logs. Option C is missing the `-c` container selector and is also malformed: `kubectl logs -p` expects the Pod name (and optionally container); `ruby` is not a flag positionally correct here. Option D has argument order incorrect and mixes Pod and container names in the wrong places.

Operationally, this is a common Kubernetes troubleshooting workflow: if a container restarts quickly, current logs may be short or empty, and the actionable crash output is in the previous instance logs. Using `kubectl logs -p` often reveals stack traces, fatal errors, or misconfiguration messages. In multicontainer Pods, always pair `-p` with `-c` to ensure you're looking at the right container.

Therefore, the verified correct answer is A.

Question: 98

A Kubernetes is an abstraction that defines a logical set of Pods and a policy by which to access them.

- A. Selector
- B. Controller
- C. Service
- D. Job

Answer: C

Explanation:

A Kubernetes Service is the abstraction that defines a logical set of Pods and the policy for accessing them, so C is correct. Pods are ephemeral: their IPs change as they are recreated, rescheduled, or scaled. A Service solves this by providing a stable endpoint (DNS name and virtual IP) and routing rules that send traffic to the current healthy Pods backing the Service.

A Service typically uses a label selector to identify which Pods belong to it. Kubernetes then maintains endpoint data (Endpoints/EndpointSlice) for those Pods and uses the cluster dataplane (kube-proxy or eBPF-based implementations) to forward traffic from the Service IP/port to one of the backend Pod IPs. This is what the question means by "logical set of Pods" and "policy by which to access them" (for example, round-robin-like distribution depending on dataplane, session affinity options, and how ports map via targetPort).

Option A (Selector) is only the query mechanism used by Services and controllers; it is not itself the access abstraction.

Option B (Controller) is too generic; controllers reconcile desired state but do not provide stable network access policies.

Option D (Job) manages run-to-completion tasks and is unrelated to network access abstraction.

Services can be exposed in different ways: ClusterIP (internal), NodePort, LoadBalancer, and ExternalName. Regardless of type, the core Service concept remains: stable access to a dynamic set of Pods. This is foundational to Kubernetes networking and microservice communication, and it is why Service discovery via DNS works effectively across rolling updates and scaling events.

Thus, the correct answer is Service (C).

Question: 99

How many hosts are required to set up a highly available Kubernetes cluster when using an external etcd topology?

- A. Four hosts. Two for control plane nodes and two for etcd nodes.
- B. Four hosts. One for a control plane node and three for etcd nodes.
- C. Three hosts. The control plane nodes and etcd nodes share the same host.
- D. Six hosts. Three for control plane nodes and three for etcd nodes.

Answer: D

Explanation:

In a highly available (HA) Kubernetes control plane using an external etcd topology, you typically run three control plane nodes and three separate etcd nodes, totaling six hosts, making D correct. HA design relies on quorum-based consensus: etcd uses Raft and requires a majority of members available to make progress. Running three etcd members is the common minimum for HA because it tolerates one member failure while maintaining quorum (2/3).

In the external etcd topology, etcd is decoupled from the control plane nodes. This separation improves fault isolation: if a control plane node fails or is replaced, etcd remains stable and independent; likewise, etcd maintenance can be handled separately. Kubernetes API servers (often multiple instances behind a load balancer) talk to the external etcd cluster for storage of cluster state.

Options A and B propose four hosts, but they break common HA/quorum best practices. Two etcd nodes do not form a robust quorum configuration (a two-member etcd cluster cannot tolerate a single failure without losing quorum).

One control plane node is not HA for the API server/scheduler/controller-manager components. Option C describes a stacked etcd topology (control plane + etcd on same hosts), which can be HA with three hosts, but the question explicitly says external etcd, not stacked. In stacked topology, you often use three control plane nodes each running an etcd member. In external topology, you use three control plane + three etcd.

Operationally, external etcd topology is often used when you want dedicated resources, separate lifecycle management, or stronger isolation for the datastore. It can reduce blast radius but increases infrastructure footprint and operational complexity (TLS, backup/restore, networking). Still, for the canonical HA external-etcd pattern, the expected answer is six hosts: 3 control plane + 3 etcd.

Question: 100

Which of these events will cause the kube-scheduler to assign a Pod to a node?

- A. When the Pod crashes because of an error.
- B. When a new node is added to the Kubernetes cluster.
- C. When the CPU load on the node becomes too high.
- D. When a new Pod is created and has no assigned node.

Answer: D

Explanation:

The kube-scheduler assigns a node to a Pod when the Pod is unscheduled—meaning it exists in the API server but has no `spec.nodeName` set. The event that triggers scheduling is therefore: a new Pod is created and has no assigned node, which is option D.

Kubernetes scheduling is declarative and event-driven. The scheduler continuously watches for Pods that are in a “Pending” unscheduled state. When it sees one, it runs a scheduling cycle: filtering nodes that cannot run the Pod (insufficient resources based on requests, taints/tolerations, node selectors/affinity rules, topology spread constraints), then scoring the remaining feasible nodes to pick the best candidate. Once selected, the scheduler “binds” the Pod to that node by updating the Pod’s `spec.nodeName`. After that, kubelet on the chosen node takes over to pull images and start containers.

Option A (Pod crashes) does not directly cause scheduling. If a container crashes, kubelet may restart it on the same node according to restart policy. If the Pod itself is replaced (e.g., by a controller like a Deployment creating a new Pod), that new Pod will be scheduled because it’s unscheduled—but the crash event itself isn’t the scheduler’s trigger. Option B (new node added) might create more capacity and affect future scheduling decisions, but it does not by itself trigger assigning a particular Pod; scheduling still happens because there are unscheduled Pods. Option C (CPU load high) is not a scheduling trigger; scheduling is based on declared requests and constraints, not instantaneous node

CPU load (that’s a common misconception).

So the correct, Kubernetes-architecture answer is D: kube-scheduler assigns nodes to Pods that are newly created (or otherwise pending) and have no assigned node.

Question: 101

What helps an organization to deliver software more securely at a higher velocity?

- A. Kubernetes
- B. apt-get
- C. Docker Images
- D. CI/CD Pipeline

Answer: D

Explanation:

A CI/CD pipeline is a core practice/tooling approach that enables organizations to deliver software faster and more securely, so D is correct. CI (Continuous Integration) automates building and testing code changes frequently, reducing integration risk and catching defects early. CD (Continuous Delivery/Deployment) automates releasing validated builds into environments using consistent, repeatable steps—reducing manual errors and enabling rapid iteration.

Security improves because automation enables standardized checks on every change: static analysis, dependency scanning, container image scanning, policy validation, and signing/verification steps can be integrated into the pipeline. Instead of relying on ad-hoc human processes, security controls become repeatable gates. In Kubernetes environments, pipelines commonly build container images, run tests, publish artifacts to registries, and then deploy via manifests, Helm, or GitOps controllers—keeping deployments consistent and auditable.

Option A (Kubernetes) is a platform that helps run and manage workloads, but by itself it doesn't guarantee secure high-velocity delivery. It provides primitives (rollouts, declarative config, RBAC),

yet the delivery workflow still needs automation. Option B (apt-get) is a package manager for Debian-based systems and is not a delivery pipeline. Option C (Docker Images) are artifacts; they improve portability and repeatability, but they don't provide the end-to-end automation of building, testing, promoting, and deploying across environments.

In cloud-native application delivery, the pipeline is the "engine" that turns code changes into safe production releases. Combined with Kubernetes' declarative deployment model (Deployments, rolling updates, health probes), a CI/CD pipeline supports frequent releases with controlled rollouts, fast rollback, and strong auditability. That is exactly what the question is targeting. Therefore, the verified answer is D.

Question: 102

Which resource do you use to attach a volume in a Pod?

- A. StorageVolume
- B. PersistentVolume
- C. StorageClass
- D. PersistentVolumeClaim

Answer: D

Explanation:

In Kubernetes, Pods typically attach persistent storage by referencing a PersistentVolumeClaim (PVC), making D correct. A PVC is a user's request for storage with specific requirements (size, access mode, storage class). Kubernetes then binds the PVC to a matching PersistentVolume (PV) (either preprovisioned statically or created dynamically via a StorageClass and CSI provisioner). The Pod does not directly attach a PV; it references the PVC, and Kubernetes handles the binding and mounting.

This design separates responsibilities: administrators (or CSI drivers) manage PV provisioning and backend storage details, while developers consume storage via PVCs. In a Pod spec, you define a volume of type persistentVolumeClaim and set claimName: <pvc-name>, then mount that volume into containers at a path. The kubelet coordinates with the CSI driver (or in-tree plugin depending on environment) to attach/mount the underlying storage to the node and then into the Pod.

Option B (PersistentVolume) is not directly referenced by Pods; PVs are cluster resources that represent actual storage. Pods don't "pick" PVs; claims do. Option C (StorageClass) defines provisioning parameters (e.g., disk type, replication, binding mode) but is not what a Pod references to mount a volume. Option A is not a Kubernetes resource type.

Operationally, using PVCs enables dynamic provisioning and portability: the same Pod spec can be deployed across clusters where the StorageClass name maps to appropriate backend storage. It also supports lifecycle controls like reclaim policies (Delete/Retain) and snapshot/restore workflows depending on CSI capabilities.

So the Kubernetes resource you use in a Pod to attach a persistent volume is PersistentVolumeClaim, option D.

Question: 103

Which key-value store is used to persist Kubernetes cluster data?

- A. etcd
- B. ZooKeeper
- C. ControlPlaneStore
- D. Redis

Answer: A

Explanation:

Kubernetes stores its cluster state (API objects) in etcd, making A correct. etcd is a distributed, strongly consistent key-value store that serves as the source of truth for the Kubernetes control plane. When you create or update objects such as Pods, Deployments, ConfigMaps, Secrets, or Nodes, the kube-apiserver validates the request and then persists the desired state into etcd.

Controllers and the scheduler watch the API for changes (which ultimately reflect etcd state) and reconcile the cluster to match that desired state.

etcd's consistency guarantees are crucial. Kubernetes relies on accurate, up-to-date state to make scheduling decisions, enforce RBAC/admission policies, coordinate leader elections, and ensure controllers behave correctly. etcd uses the Raft consensus algorithm to replicate data among members and requires quorum for writes, enabling fault tolerance when deployed in HA configurations (commonly three or five members).

The other options are incorrect in Kubernetes' standard architecture. ZooKeeper is a distributed coordination system used by some other platforms, but Kubernetes does not use it as its primary datastore. Redis is an in-memory data store used for caching or messaging, not as Kubernetes' authoritative state store. "ControlPlaneStore" is not a standard Kubernetes component.

Operationally, etcd health is one of the most important determinants of cluster reliability. Slow disk I/O or unstable networking can degrade etcd performance and cause API latency spikes. Backup and restore procedures for etcd are critical disaster-recovery practices, and securing etcd (TLS, access restrictions) is essential because it may contain sensitive data (e.g., Secrets—often base64-encoded, and optionally encrypted at rest depending on configuration).

Therefore, the verified Kubernetes datastore is etcd, option A.

Question: 104

What Linux namespace is shared by default by containers running within a Kubernetes Pod?

- A. Host Network
- B. Network
- C. Process ID
- D. Process Name

Answer: B

Explanation:

By default, containers in the same Kubernetes Pod share the network namespace, which means they share the same IP address and port space. Therefore, the correct answer is B (Network).

This shared network namespace is a key part of the Pod abstraction. Because all containers in a Pod share networking, they can communicate with each other over localhost and coordinate tightly, which is the basis for patterns like sidecars (service mesh proxies, log shippers, config reloaders). It also means containers must coordinate port usage: if two containers try to bind the same port on 0.0.0.0, they'll conflict because they share the same port namespace.

Option A ("Host Network") is different: `hostNetwork: true` is an optional Pod setting that puts the Pod into the node's network namespace, not the Pod's shared namespace. It is not the default and is generally used sparingly due to security and port-collision risks. Option C ("Process ID") is not shared by default in Kubernetes; PID namespace sharing requires explicitly enabling process namespace sharing (e.g., `shareProcessNamespace: true`). Option D ("Process Name") is not a Linux namespace concept.

The Pod model also commonly implies shared storage volumes (if defined) and shared IPC namespace in some configurations, but the universally shared-by-default namespace across containers in the same Pod is the network namespace. This default behavior is why Kubernetes documentation explains a Pod as a "logical host" for one or more containers: the containers are colocated and share certain namespaces as if they ran on the same host.

So, the correct, verified answer is B: containers in the same Pod share the Network namespace by default.

Question: 105

What is a Dockerfile?

- A. A bash script that is used to automatically build a docker image.
- B. A config file that defines which image registry a container should be pushed to.
- C. A text file that contains all the commands a user could call on the command line to assemble an image.
- D. An image layer created by a running container stored on the host.

Answer: C

Explanation:

A Dockerfile is a text file that contains a sequence of instructions used to build a container image, so C is correct. These instructions include choosing a base image (FROM), copying files (COPY/ADD), installing dependencies (RUN), setting environment variables (ENV), defining working directories (WORKDIR), exposing ports (EXPOSE), and specifying the default startup command (CMD/ENTRYPOINT). When you run docker build (or compatible tools like BuildKit), the builder executes these instructions to produce an image composed of immutable layers.

In cloud-native application delivery, Dockerfiles (more generally, OCI image build definitions) are a key step in the supply chain. The resulting image artifact is what Kubernetes runs in Pods. Best practices include using minimal base images, pinning versions, avoiding embedding secrets, and using multi-stage builds to keep runtime images small. These practices improve security and performance, and make delivery pipelines more reliable.

Option A is incorrect because a Dockerfile is not a bash script, even though it can run shell commands through RUN. Option B is incorrect because registry destinations are handled by tooling and tagging/push commands (or CI pipeline configuration), not by the Dockerfile itself. Option D is incorrect because an image layer created by a running container is more closely related to container filesystem changes and commits; a Dockerfile is the build recipe, not a runtime-generated layer.

Although the question uses "Dockerfile," the concept maps well to OCI-based container image creation generally: you define a reproducible build recipe that produces an immutable image artifact. That artifact is then versioned, scanned, signed, stored in a registry, and deployed to Kubernetes through manifests/Helm/GitOps. Therefore, C is the correct and verified definition.

Question: 106

What does the "nodeSelector" within a PodSpec use to place Pods on the target nodes?

- A. Annotations
- B. IP Addresses
- C. Hostnames
- D. Labels

Answer: D

Explanation:

nodeSelector is a simple scheduling constraint that matches node labels, so the correct answer is D (Labels). In Kubernetes, nodes have key/value labels (for example, disktype=ssd, topology.kubernetes.io/zone=us-east-1a, kubernetes.io/os=linux). When you set spec.nodeSelector in a Pod template, you provide a map of required label key/value pairs. The kube-scheduler will then only consider nodes that have all those labels with matching values as eligible placement targets for that Pod.

This is different from annotations: annotations are also key/value metadata, but they are not intended for selection logic and are not used by the scheduler for nodeSelector. IP addresses and hostnames are not the mechanism used by nodeSelector either. While Kubernetes nodes do have hostnames and IPs, nodeSelector specifically operates on labels because labels are designed for selection, grouping, and placement constraints.

Operationally, nodeSelector is the most basic form of node placement control. It is commonly used to pin workloads to specialized hardware (GPU nodes), compliance zones, or certain OS/architecture pools. However, it has limitations: it only supports exact match on labels and cannot express more complex rules (like "in this set of zones" or "prefer but don't require"). For that, Kubernetes offers node affinity (requiredDuringSchedulingIgnoredDuringExecution, preferredDuringSchedulingIgnoredDuringExecution) which supports richer expressions.

Still, the underlying mechanism is the same concept: the scheduler evaluates your Pod's placement requirements against node metadata, and for nodeSelector, that metadata is labels. Therefore, the verified correct answer is D.

Question: 107

What do Deployments and StatefulSets have in common?

- A. They manage Pods that are based on an identical container spec.
- B. They support the OnDelete update strategy.
- C. They support an ordered, graceful deployment and scaling.
- D. They maintain a sticky identity for each of their Pods.

Answer: A

Explanation:

Both Deployments and StatefulSets are Kubernetes workload controllers that manage a set of Pods created from a Pod template, meaning they manage Pods based on an identical container specification (a shared Pod template). That is why A is correct. In both cases, you declare a desired state (replicas, container images, environment variables, volumes, probes, etc.) in spec.template, and the controller ensures the cluster converges toward that state by creating, updating, or replacing Pods.

The differences are what make the other options incorrect. OnDelete update strategy is associated with StatefulSets (it's one of their update strategies), but it is not a shared, defining behavior across both controllers, so B is not "in common." Ordered, graceful deployment and scaling is a hallmark of StatefulSets (ordered pod creation/termination and stable identities) rather than Deployments, so C is not shared. Sticky identity per Pod (stable network identity and stable storage identity per replica, commonly via StatefulSet + headless Service) is specifically a StatefulSet characteristic, not a Deployment feature, so D is not common.

A useful way to think about it is: both controllers manage replicas of a Pod template, but they differ in semantics.

Deployments are designed primarily for stateless workloads and typically focus on rolling updates and scalable replicas where any instance is interchangeable. StatefulSets are designed for stateful workloads and add identity and ordering guarantees: each replica gets a stable name (like db-0, db-1) and often stable PersistentVolumeClaims.

So the shared commonality the question is testing is the basic workload-controller pattern: both controllers manage Pods created from a common template (identical container spec). Therefore, A is the verified answer.

Question: 108

What is the practice of bringing financial accountability to the variable spend model of cloud resources?

- A. FaaS
- B. DevOps
- C. CloudCost
- D. FinOps

Answer: D

Explanation:

The practice of bringing financial accountability to cloud spending—where costs are variable and usage-based—is called FinOps, so D is correct. FinOps (Financial Operations) is an operating model and culture that helps organizations manage cloud costs by connecting engineering, finance, and business teams. Because cloud resources can be provisioned quickly and billed dynamically, traditional budgeting approaches often fail to keep pace. FinOps addresses this by introducing shared visibility, governance, and optimization processes that enable teams to make cost-aware decisions while still moving fast.

In Kubernetes and cloud-native architectures, variable spend shows up in many ways: autoscaling node pools, over-provisioned resource requests, idle clusters, persistent volumes, load balancers, egress traffic, managed services, and observability tooling. FinOps practices encourage tagging/labeling for cost attribution, defining cost KPIs, enforcing budget guardrails, and continuously optimizing usage (right-sizing resources, scaling policies, turning off unused environments, and selecting cost-effective architectures).

Why the other options are incorrect: FaaS (Function as a Service) is a compute model (serverless), not a financial accountability practice. DevOps is a cultural and technical practice focused on collaboration and delivery speed, not specifically cloud cost accountability (though it can complement FinOps). CloudCost is not a widely recognized standard term in the way FinOps is.

In practice, FinOps for Kubernetes often involves improving resource efficiency: aligning requests/limits with real usage, using HPA/VPA appropriately, selecting instance types that match workload profiles, managing cluster autoscaler settings, and allocating shared platform costs to

teams via labels/namespaces. It also includes forecasting and anomaly detection, because cloudnative spend can spike quickly due to misconfigurations (e.g., runaway autoscaling or excessive log ingestion).

So, the correct term for financial accountability in cloud variable spend is FinOps (D).

Question: 109

What is a best practice to minimize the container image size?

- A. Use a DockerFile.
- B. Use multistage builds.
- C. Build images with different tags.
- D. Add a build.sh script.

Answer: B

Explanation:

A proven best practice for minimizing container image size is to use multi-stage builds, so B is correct. Multi-stage builds allow you to separate the “build environment” from the “runtime environment.” In the first stage, you can use a full-featured base image (with compilers, package managers, and build tools) to compile your application or assemble artifacts. In the final stage, you copy only the resulting binaries or necessary runtime assets into a much smaller base image (for example, a distroless image or a slim OS image). This dramatically reduces the final image size because it excludes compilers, caches, and build dependencies that are not needed at runtime.

In cloud-native application delivery, smaller images matter for several reasons. They pull faster, which speeds up deployments, rollouts, and scaling events (Pods become Ready sooner). They also reduce attack surface by removing unnecessary packages, which helps security posture and scanning results. Smaller images tend to be simpler and more reproducible, improving reliability across environments.

Option A is not a size-minimization practice: using a Dockerfile is simply the standard way to define how to build an image; it doesn't inherently reduce size. Option C (different tags) changes image identification but not size. Option D (a build script) may help automation, but it doesn't guarantee smaller images; the image contents are determined by what ends up in the layers.

Multi-stage builds are commonly paired with other best practices: choosing minimal base images, cleaning package caches, avoiding copying unnecessary files (use `.dockerignore`), and reducing layer churn. But among the options, the clearest and most directly correct technique is multi-stage builds.

Therefore, the verified answer is B.

Question: 110

Which tools enable Kubernetes HorizontalPodAutoscalers to use custom, application-generated metrics to trigger scaling events?

- A. Prometheus and the prometheus-adapter.
- B. Graylog and graylog-autoscaler metrics.
- C. Graylog and the kubernetes-adapter.
- D. Grafana and Prometheus.

Answer: A

Explanation:

To scale on custom, application-generated metrics, the Horizontal Pod Autoscaler (HPA) needs those metrics exposed through the Kubernetes custom metrics (or external metrics) API. A common and Kubernetes-documented approach is Prometheus + prometheus-adapter, making A correct. Prometheus scrapes application metrics (for example, request rate, queue depth, in-flight requests) from /metrics endpoints. The prometheus-adapter then translates selected Prometheus time series into the Kubernetes Custom Metrics API so the HPA controller can fetch them and make scaling decisions.

Why not the other options: Grafana is a visualization tool; it does not provide the metrics API translation layer required by HPA, so "Grafana and Prometheus" is incomplete. Graylog is primarily a log management system; it's not the standard solution for feeding custom metrics into HPA via the Kubernetes metrics APIs. The "kubernetes-adapter" term in option C is not the standard named adapter used in the common Kubernetes ecosystem for Prometheus-backed custom metrics (the recognized component is prometheus-adapter).

This matters operationally because HPA is not limited to CPU/memory. CPU and memory use resource metrics (often from metrics-server), but modern autoscaling often needs application signals: message queue length, requests per second, latency, or business metrics. With Prometheus and prometheus-adapter, you can define HPA rules such as "scale to maintain queue depth under X" or "scale based on requests per second per pod." This can produce better scaling behavior than CPU-based scaling alone, especially for I/O-bound services or workloads with uneven CPU profiles.

So the correct tooling combination in the provided choices is Prometheus and the prometheus-adapter, option A.

Question: 111

Which of the following is a valid PromQL query?

- A. `SELECT * from http_requests_total WHERE job=apiserver`
- B. `http_requests_total WHERE (job="apiserver")`
- C. `SELECT * from http_requests_total`
- D. `http_requests_total(job="apiserver")`

Answer: D

Explanation:

Prometheus Query Language (PromQL) uses a function-and-selector syntax, not SQL. A valid query typically starts with a metric name and optionally includes label matchers in curly braces. In the simplified quiz syntax given, the valid PromQL-style selector is best represented by D:

`http_requests_total(job="apiserver")`, so D is correct.

Conceptually, what this query means is “select time series for the metric `http_requests_total` where the job label equals `apiserver`.” In standard PromQL formatting you most often see this as: `http_requests_total{job="apiserver"}`. Many training questions abbreviate braces and focus on the idea of filtering by labels; the key is that PromQL uses label matchers rather than SQL WHERE clauses.

Options A and C are invalid because they use SQL (`SELECT * FROM ...`) which is not PromQL. Option B is also invalid because PromQL does not use the keyword WHERE. PromQL filtering is done by applying label matchers directly to the metric selector.

In Kubernetes observability, PromQL is central to building dashboards and alerts from cluster metrics. For example, you might compute rates from counters: `rate(http_requests_total{job="apiserver"}[5m])`, aggregate by labels: `sum by (code) (...)`, or alert on error ratios. Understanding the selector and label-matcher model is foundational because Prometheus metrics are multi-dimensional—labels define the slices you can filter and aggregate on.

So, within the provided options, D is the only one that follows PromQL’s metric+label-filter style and therefore is the verified correct answer.

Question: 112

Which of the following best describes horizontally scaling an application deployment?

- A. The act of adding/removing node instances to the cluster to meet demand.
- B. The act of adding/removing applications to meet demand.
- C. The act of adding/removing application instances of the same application to meet demand.
- D. The act of adding/removing resources to application instances to meet demand.

Answer: C

Explanation:

Horizontal scaling means changing how many instances of an application are running, not changing how big each instance is. Therefore, the best description is C: adding/removing application instances of the same application to meet demand.

In Kubernetes, “instances” typically correspond to Pod replicas managed by a controller like a Deployment. When you scale horizontally, you increase or decrease the replica count, which increases or decreases total throughput and resilience by distributing load across more Pods.

Option A is about cluster/node scaling (adding or removing nodes), which is infrastructure scaling typically handled by a cluster autoscaler in cloud environments. Node scaling can enable more Pods to be scheduled, but it’s not the definition of horizontal application scaling itself. Option D describes vertical scaling—adding/removing CPU or memory resources to a given instance (Pod/container) by changing requests/limits or using VPA. Option B is vague and not the standard definition.

Horizontal scaling is a core cloud-native pattern because it improves availability and elasticity. If one Pod fails, other replicas continue serving traffic. In Kubernetes, scaling can be manual (`kubectl scale deployment ... --replicas=N`) or automatic using the Horizontal Pod Autoscaler (HPA). HPA adjusts replicas based on observed metrics like CPU utilization, memory, or custom/external metrics (for example, request rate or queue length). This creates responsive systems that can handle variable traffic.

From an architecture perspective, designing for horizontal scaling often means ensuring your application is stateless (or manages state externally), uses idempotent request handling, and supports multiple concurrent instances. Stateful workloads can also scale horizontally, but usually with additional constraints (StatefulSets, sharding, quorum membership, stable identity).

So the verified definition and correct choice is C.

Question: 113

How many different Kubernetes service types can you define?

- A. 2
- B. 3
- C. 4
- D. 5

Answer: C

Explanation:

Kubernetes defines four primary Service types, which is why C (4) is correct. The commonly recognized Service spec.type values are:

ClusterIP: The default type. Exposes the Service on an internal virtual IP reachable only within the cluster. This supports typical east-west traffic between workloads.

NodePort: Exposes the Service on a static port on each node. Traffic to <NodeIP>:<NodePort> is forwarded to the Service endpoints. This is often used for simple external access in environments without load balancers, or as a building block for other systems.

LoadBalancer: Integrates with a cloud provider (or load balancer implementation) to provision an external load balancer and route traffic to the Service. This is common in managed Kubernetes.

ExternalName: Maps the Service name to an external DNS name via a CNAME record, allowing incluster clients to use a consistent Service DNS name to reach an external dependency.

Some people also talk about "Headless Services," but headless is not a separate type; it's a behavior achieved by setting clusterIP: None. Headless Services still use the Service API object but change DNS and virtual-IP behavior to return endpoint IPs directly rather than a ClusterIP. That's why the canonical count of "Service types" is four.

This question tests understanding of the Service abstraction: Service type controls how a stable service identity is exposed (internal VIP, node port, external LB, or DNS alias), while selectors/endpoints control where traffic goes (the backend Pods). Different environments will favor different types: ClusterIP for internal microservices, LoadBalancer for external exposure in cloud, NodePort for bare-metal or simple access, ExternalName for bridging to outside services.

Therefore, the verified answer is C (4).

Question: 114

What is the difference between a Deployment and a ReplicaSet?

- A. With a Deployment, you can't control the number of pod replicas.
- B. A ReplicaSet does not guarantee a stable set of replica pods running.
- C. A Deployment is basically the same as a ReplicaSet with annotations.
- D. A Deployment is a higher-level concept that manages ReplicaSets.

Answer: D

Explanation:

A Deployment is a higher-level controller that manages ReplicaSets and provides rollout/rollback behavior, so D is correct. A ReplicaSet's primary job is to ensure that a specified number of Pod replicas are running at any time, based on a label selector and Pod template. It's a fundamental "keep N Pods alive" controller.

Deployments build on that by managing the lifecycle of ReplicaSets over time. When you update a Deployment (for example, changing the container image tag or environment variables), Kubernetes creates a new ReplicaSet for the new Pod template and gradually shifts replicas from the old ReplicaSet to the new one according to the rollout strategy (RollingUpdate by default). Deployments also retain revision history, making it possible to roll back to a previous ReplicaSet if a rollout fails.

Why the other options are incorrect:

A is false: Deployments absolutely control the number of replicas via `spec.replicas` and can also be controlled by HPA.

B is false: ReplicaSets do guarantee that a stable number of replicas is running (that is their core purpose).

C is false: a Deployment is not "a ReplicaSet with annotations." It is a distinct API resource with additional controller logic for declarative updates, rollouts, and revision tracking.

Operationally, most teams create Deployments rather than ReplicaSets directly because Deployments are safer and more feature-complete for application delivery. ReplicaSets still appear in real clusters because Deployments create them automatically; you'll commonly see multiple ReplicaSets during rollout transitions. Understanding the hierarchy is crucial for troubleshooting: if Pods aren't behaving as expected, you often trace from Deployment → ReplicaSet → Pod, checking selectors, events, and rollout status.

So the key difference is: ReplicaSet maintains replica count; Deployment manages ReplicaSets and orchestrates updates. Therefore, D is the verified answer.

Question: 115

The Container Runtime Interface (CRI) defines the protocol for the communication between:

- A. The kubelet and the container runtime.
- B. The container runtime and etcd.
- C. The kube-apiserver and the kubelet.
- D. The container runtime and the image registry.

Answer: A

Explanation:

The CRI (Container Runtime Interface) defines how the kubelet talks to the container runtime, so A is correct. The kubelet is the node agent responsible for ensuring containers are running in Pods on that node. It needs a standardized way to request operations such as: create a Pod sandbox, pull an image, start/stop containers, execute commands, attach streams, and retrieve logs. CRI provides that contract so kubelet does not need runtime-specific integrations.

This interface is a key part of Kubernetes' modular design. Different container runtimes implement the CRI, allowing Kubernetes to run with containerd, CRI-O, and other CRI-compliant runtimes. This separation of concerns lets Kubernetes focus on orchestration, while runtimes focus on executing containers according to the OCI runtime spec, managing images, and handling low-level container lifecycle.

Why the other options are incorrect:

etcd is the control plane datastore; container runtimes do not communicate with etcd via CRI.

kube-apiserver and kubelet communicate using Kubernetes APIs, but CRI is not their protocol; CRI is specifically kubelet ↔ runtime.

container runtime and image registry communicate using registry protocols (image pull/push APIs),

but that is not CRI. CRI may trigger image pulls via runtime requests, yet the actual registry communication is

separate.

Operationally, this distinction matters when debugging node issues. If Pods are stuck in "ContainerCreating" due to image pull failures or runtime errors, you often investigate kubelet logs and the runtime (containerd/CRI-O) logs. Kubernetes administrators also care about CRI streaming (exec/attach/logs streaming), runtime configuration, and compatibility across Kubernetes versions.

So, the verified answer is A: the kubelet and the container runtime.

Question: 116

Which authorization-mode allows granular control over the operations that different entities can perform on different objects in a Kubernetes cluster?

- A. Webhook Mode Authorization Control
- B. Role Based Access Control
- C. Node Authorization Access Control
- D. Attribute Based Access Control

Answer: B

Explanation:

Role Based Access Control (RBAC) is the standard Kubernetes authorization mode that provides granular control over what users and service accounts can do to which resources, so B is correct. RBAC works by defining Roles (namespaced) and ClusterRoles (cluster-wide) that contain sets of rules. Each rule specifies API groups, resource types, resource names (optional), and allowed verbs such as get, list, watch, create, update, patch, and delete. You then attach these roles to identities using RoleBindings or ClusterRoleBindings.

This gives fine-grained, auditable access control. For example, you can allow a CI service account to create and patch Deployments only in a specific namespace, while restricting it from reading Secrets.

You can allow developers to view Pods and logs but prevent them from changing cluster-wide networking resources. This

is exactly the “granular control over operations on objects” described by the question.

Why other options are not the best answer: “Webhook mode” is an authorization mechanism where Kubernetes calls an external service to decide authorization. While it can be granular depending on the external system, Kubernetes’ common built-in answer for granular object-level control is RBAC. “Node authorization” is a specialized authorizer for kubelets/nodes to access resources they need; it’s not the general-purpose system for all cluster entities. ABAC (Attribute-Based Access Control) is an older mechanism and is not the primary recommended authorization model; it can be expressive but is less commonly used and not the default best-practice for Kubernetes authorization today.

In Kubernetes security practice, RBAC is typically paired with authentication (certs/OIDC), admission controls, and namespaces to build a defense-in-depth security posture. RBAC policy is also central to least privilege: granting only what is necessary for a workload or user role to function. This reduces blast radius if credentials are compromised.

Therefore, the verified answer is B: Role Based Access Control.

Question: 117

Which of the following is a correct definition of a Helm chart?

- A. A Helm chart is a collection of YAML files bundled in a tar.gz file and can be applied without decompressing it.
- B. A Helm chart is a collection of JSON files and contains all the resource definitions to run an application on Kubernetes.
- C. A Helm chart is a collection of YAML files that can be applied on Kubernetes by using the kubectl tool.
- D. A Helm chart is similar to a package and contains all the resource definitions to run an application on Kubernetes.

Answer: D

Explanation:

A Helm chart is best described as a package for Kubernetes applications, containing the resource definitions (as templates) and metadata needed to install and manage an application—so D is correct. Helm is a package manager for Kubernetes; the chart is the packaging format. Charts include a Chart.yaml (metadata), a values.yaml (default configuration values), and a templates/ directory containing Kubernetes manifests written as templates. When you install a chart, Helm renders those templates into concrete Kubernetes YAML manifests by substituting values, then applies them to the cluster.

Option A is misleading/incomplete. While charts are often distributed as a compressed tarball (.tgz), the defining feature is not “YAML bundled in tar.gz” but the packaging and templating model that supports install/upgrade/rollback. Option B is incorrect because Helm charts are not “collections of JSON files” by definition; Kubernetes resources can be expressed as YAML or JSON, but Helm charts overwhelmingly use templated YAML. Option C is incorrect because charts are not simply YAML applied by kubectl; Helm manages releases, tracks installed resources, and supports upgrades and rollbacks. Helm uses Kubernetes APIs under the hood, but the value of Helm is the lifecycle and packaging system, not “kubectl apply.”

In cloud-native application delivery, Helm helps standardize deployments across environments (dev/stage/prod) by externalizing configuration through values. It reduces copy/paste and supports reuse via dependencies and subcharts. Helm also supports versioning of application packages, allowing teams to upgrade predictably and roll back if needed—critical for production change management.

So, the correct and verified definition is D: a Helm chart is like a package containing the resource definitions needed to run an application on Kubernetes.

Question: 118

Which of the following sentences is true about namespaces in Kubernetes?

- A. You can create a namespace within another namespace in Kubernetes.
- B. You can create two resources of the same kind and name in a namespace.
- C. The default namespace exists when a new cluster is created.
- D. All the objects in the cluster are namespaced by default.

Answer: C

Explanation:

The true statement is C: the default namespace exists when a new cluster is created. Namespaces are a Kubernetes mechanism for partitioning cluster resources into logical groups. When you set up a cluster, Kubernetes creates some initial namespaces (including default, and commonly kube-system, kube-public, and kube-node-lease). The default namespace is where resources go if you don't specify a namespace explicitly.

Option A is false because namespaces are not hierarchical; Kubernetes does not support “namespaces inside namespaces.” Option B is false because within a given namespace, resource names must be unique per resource kind. You can't have two Deployments with the same name in the same namespace. You can have a Deployment named web in one namespace and another Deployment named web in a different namespace—namespaces provide that scope boundary. Option D is false because not all objects are namespaced. Many resources are cluster-scoped (for example, Nodes, PersistentVolumes, ClusterRoles, ClusterRoleBindings, and StorageClasses). Namespaces apply only to namespaced resources.

Operationally, namespaces support multi-tenancy and environment separation (dev/test/prod), RBAC scoping, resource quotas, and policy boundaries. For example, you can grant a team access only to their namespace and enforce quotas that prevent them from consuming excessive CPU/memory. Namespaces also make organization and cleanup easier: deleting a namespace removes most namespaced resources inside it (subject to finalizers).

So, the verified correct statement is C: the default namespace exists upon cluster creation.

Question: 119

How does Horizontal Pod autoscaling work in Kubernetes?

- A. The Horizontal Pod Autoscaler controller adds more CPU or memory to the pods when the load is above the configured threshold, and reduces CPU or memory when the load is below.
- B. The Horizontal Pod Autoscaler controller adds more pods when the load is above the configured threshold, but does not reduce the number of pods when the load is below.
- C. The Horizontal Pod Autoscaler controller adds more pods to the specified DaemonSet when the load is above the configured threshold, and reduces the number of pods when the load is below.
- D. The Horizontal Pod Autoscaler controller adds more pods when the load is above the configured threshold, and reduces the number of pods when the load is below.

Answer: D

Explanation:

Horizontal Pod Autoscaling (HPA) adjusts the number of Pod replicas for a workload controller (most commonly a Deployment) based on observed metrics, increasing replicas when load is high and decreasing when load drops.

That matches D, so D is correct.

HPA does not add CPU or memory to existing Pods—that would be vertical scaling (VPA). Instead, HPA changes spec.replicas on the target resource, and the controller then creates or removes Pods accordingly. HPA commonly scales based on CPU utilization and memory (resource metrics), and it can also scale using custom or external metrics if those are exposed via the appropriate Kubernetes metrics APIs.

Option A is vertical scaling behavior, not HPA. Option B is incorrect because HPA can scale down as well as up (subject to stabilization windows and configuration), so it's not "scale up only." Option C is incorrect because HPA does not scale DaemonSets in the usual model; DaemonSets are designed to run one Pod per node (or per selected nodes) rather than a replica count. HPA targets resources like Deployments, ReplicaSets (via Deployment), and StatefulSets in typical usage, where replica count is a meaningful knob.

Operationally, HPA works as a control loop: it periodically reads metrics (for example, via metricsserver for CPU/memory, or via adapters for custom metrics), compares the current value to the desired target, and calculates a desired replica count within min/max bounds. To avoid flapping, HPA includes stabilization behavior and cooldown logic so it doesn't scale too aggressively in response to short spikes or dips. You can configure minimum and maximum replicas and behavior policies to tune responsiveness.

In cloud-native systems, HPA is a key elasticity mechanism: it enables services to handle variable traffic while controlling cost by scaling down during low demand. Therefore, the verified correct answer is D.

Question: 120

What is a Pod?

- A. A networked application within Kubernetes.
- B. A storage volume within Kubernetes.
- C. A single container within Kubernetes.
- D. A group of one or more containers within Kubernetes.

Answer: D

Explanation:

A Pod is the smallest deployable/schedulable unit in Kubernetes and consists of a group of one or more containers that are deployed together on the same node—so D is correct. The key idea is that Kubernetes schedules Pods, not individual containers. Containers in the same Pod share important runtime context: they share the same network namespace (one

Pod IP and port space) and can share storage volumes defined at the Pod level. This is why a Pod is often described as a “logical host” for its containers.

Most Pods run a single container, but multi-container Pods are common for sidecar patterns. For example, an application container might run alongside a service mesh proxy sidecar, a log shipper, or a config reloader. Because these containers share localhost networking, they can communicate efficiently without exposing extra network endpoints. Because they can share volumes, one container can produce files that another consumes (for example, writing logs to a shared volume).

Options A and B are incorrect because a Pod is not “an application” abstraction nor is it a storage volume. Pods can host applications, but they are the execution unit for containers rather than the application concept itself. Option C is incorrect because a Pod is not limited to a single container; “one or more containers” is fundamental to the Pod definition.

Operationally, understanding Pods is essential because many Kubernetes behaviors key off Pods: Services select Pods (typically by labels), autoscalers scale Pods (replica counts), probes determine

Pod readiness/liveness, and scheduling constraints place Pods on nodes. When a Pod is replaced (for example during a Deployment rollout), a new Pod is created with a new UID and potentially a new IP—reinforcing why Services exist to provide stable access.

Therefore, the verified correct answer is D: a Pod is a group of one or more containers within Kubernetes.

Question: 121

What element allows Kubernetes to run Pods across the fleet of nodes?

- A. The node server.
- B. The etcd static pods.
- C. The API server.
- D. The kubelet.

Answer: D

Explanation:

The correct answer is D (the kubelet) because the kubelet is the node agent responsible for actually running Pods on each node. Kubernetes can orchestrate workloads across many nodes because every worker node (and control-plane node that runs workloads) runs a kubelet that continuously watches the API server for PodSpecs assigned to that node and then ensures the containers described by those PodSpecs are started and kept running. In other words, the kube-scheduler decides where a Pod should run (sets spec.nodeName), but the kubelet is what makes the Pod run on that chosen node.

The kubelet integrates with the container runtime (via CRI) to pull images, create sandboxes, start containers, and manage their lifecycle. It also reports node and Pod status back to the control plane, executes liveness/readiness/startup probes, mounts volumes, and performs local housekeeping that keeps the node aligned with the declared desired state.

This node-level reconciliation loop is a key Kubernetes pattern: the control plane declares intent, and the kubelet enforces it on the node.

Option C (API server) is critical but does not run Pods; it is the control plane's front door for storing and serving cluster state. Option A ("node server") is not a Kubernetes component. Option B (etcd static pods) is a misunderstanding: etcd is the datastore for Kubernetes state and may run as static Pods in some installations, but it is not the mechanism that runs user workloads across nodes.

So, Kubernetes runs Pods "across the fleet" because each node has a kubelet that can realize scheduled PodSpecs locally and keep them healthy over time.

Question: 122

What is the Kubernetes object used for running a recurring workload?

- A. Job
- B. Batch
- C. DaemonSet
- D. CronJob

Answer: D

Explanation:

A recurring workload in Kubernetes is implemented with a CronJob, so the correct choice is D. A CronJob is a controller that creates Jobs on a schedule defined in standard cron format (minute, hour, day of month, month, day of week). This makes CronJobs ideal for periodic tasks like backups, report generation, log rotation, and cleanup tasks.

A Job (option A) is run-to-completion but is typically a one-time execution; it ensures that a specified number of Pods successfully terminate. You can use a Job repeatedly, but something else must create it each time—CronJob is that built-in scheduler. Option B (“Batch”) is not a standard workload resource type (batch is an API group, not the object name used here). Option C (DaemonSet) ensures one Pod runs on every node (or selected nodes), which is not “recurring,” it’s “always present per node.”

CronJobs include operational controls that matter in real clusters. For example, concurrencyPolicy controls what happens if a scheduled run overlaps with a previous run (Allow, Forbid, Replace). startingDeadlineSeconds can handle missed schedules (e.g., if the controller was down). History limits (successfulJobsHistoryLimit, failedJobsHistoryLimit) help manage cleanup and troubleshooting. Each scheduled execution results in a Job with its own Pods, which can be inspected with `kubectl get jobs` and `kubectl logs`.

So the correct Kubernetes object for a recurring workload is CronJob (D): it provides native scheduling and creates Jobs automatically according to the defined cadence.

Question: 123

In the DevOps framework and culture, who builds, automates, and offers continuous delivery tools for developer teams?

- A. Application Users
- B. Application Developers
- C. Platform Engineers
- D. Cluster Operators

Answer: C

Explanation:

The correct answer is C (Platform Engineers). In modern DevOps and platform operating models, platform engineering

teams build and maintain the shared delivery capabilities that product/application teams use to ship software safely and quickly. This includes CI/CD pipeline templates, standardized build and test automation, artifact management (registries), deployment tooling (Helm/Kustomize/GitOps), secrets management patterns, policy guardrails, and paved-road workflows that reduce cognitive load for developers.

While application developers (B) write the application code and often contribute pipeline steps for their service, the “build, automate, and offer tooling for developer teams” responsibility maps directly to platform engineering: they provide the internal platform that turns Kubernetes and cloud services into a consumable product. This is especially common in Kubernetes-based organizations where you want consistent deployment standards, repeatable security checks, and uniform observability.

Cluster operators (D) typically focus on the health and lifecycle of the Kubernetes clusters themselves: upgrades, node pools, networking, storage, cluster security posture, and control plane reliability. They may work closely with platform engineers, but “continuous delivery tools for developer teams” is broader than cluster operations. Application users (A) are consumers of the software, not builders of delivery tooling.

In cloud-native application delivery, this division of labor is important: platform engineers enable higher velocity with safety by automating the software supply chain—builds, tests, scans, deploys, progressive delivery, and rollback. Kubernetes provides the runtime substrate, but the platform team makes it easy and safe for developers to use it repeatedly and consistently across many services.

Therefore, Platform Engineers (C) is the verified correct choice.

Question: 124

Which kubectl command is useful for collecting information about any type of resource that is active in a Kubernetes cluster?

- A. describe
- B. list
- C. expose
- D. explain

Answer: A

Explanation:

The correct answer is A (describe), used as `kubectl describe <resource> <name>`. `kubectl describe` is a troubleshooting-focused command that provides a rich, human-readable view of a specific live object in the cluster, including key fields, status, and—crucially—Events related to that object. This makes it extremely useful for “collecting information” about almost any active resource: Pods, Deployments, Nodes, Services, PersistentVolumeClaims, and more.

`kubectl get` (not listed) is typically used for listing objects and their summary fields, but `kubectl describe` goes deeper: for a Pod it will show container images, resource requests/limits, probes, mounted volumes, node assignment, IPs, conditions, and recent scheduling/pulling/starting events. For a Node it shows capacity/allocatable resources, labels/taints, conditions, and node events. Those event details often explain why something is Pending, failing to pull images, failing readiness checks, or being evicted.

Option B (“list”) is not a standard `kubectl` subcommand for retrieving resource information (you would use `get` for listing). Option C (expose) is for creating a Service to expose a resource (like a Deployment). Option D (explain) is for viewing API schema/field documentation (e.g., `kubectl explain deployment.spec.replicas`) and does not report what is currently happening in the cluster.

So, for gathering detailed live diagnostics about a resource in the cluster, the best `kubectl` command is `kubectl describe`, which corresponds to option A.

Question: 125

The cloud native architecture centered around microservices provides a strong system that ensures

- A. fallback
- B. resiliency
- C. failover
- D. high reachability

Answer: B

Explanation:

The best answer is B (resiliency). A microservices-centered cloud-native architecture is designed to build systems that continue to operate effectively under change and failure. "Resiliency" is the umbrella concept: the ability to tolerate faults, recover from disruptions, and maintain acceptable service levels through redundancy, isolation, and automated recovery.

Microservices help resiliency by reducing blast radius. Instead of one monolith where a single defect can take down the entire application, microservices separate concerns into independently deployable components. Combined with Kubernetes, you get resiliency mechanisms such as replication (multiple Pod replicas), self-healing (restart and reschedule on failure), rolling updates, health probes, and service discovery/load balancing. These enable the platform to detect and replace failing instances automatically, and to keep traffic flowing to healthy backends.

Options C (failover) and A (fallback) are resiliency techniques but are narrower terms. Failover usually refers to switching to a standby component when a primary fails; fallback often refers to degraded behavior (cached responses, reduced features). Both can exist in microservice systems, but the broader architectural guarantee microservices aim to support is resiliency overall. Option D ("high reachability") is not the standard term used in cloud-native design and doesn't capture the intent as precisely as resiliency.

In practice, achieving resiliency also requires good observability and disciplined delivery: monitoring/alerts, tracing across service boundaries, circuit breakers/timeouts/retries, and progressive delivery patterns. Kubernetes provides platform primitives, but resilient microservices also need careful API design and failure-mode thinking.

So the intended and verified completion is resiliency, option B.

Question: 126

Which of the following is the correct command to run an nginx deployment with 2 replicas?

- A. `kubectl run deploy nginx --image=nginx --replicas=2`
- B. `kubectl create deploy nginx --image=nginx --replicas=2`
- C. `kubectl create nginx deployment --image=nginx --replicas=2`
- D. `kubectl create deploy nginx --image=nginx --count=2`

Answer: B

Explanation:

The correct answer is B: `kubectl create deploy nginx --image=nginx --replicas=2`. This uses `kubectl create deployment` (shorthand `create deploy`) to generate a Deployment resource named `nginx` with the specified container image. The `--replicas=2` flag sets the desired replica count, so Kubernetes will create two Pod replicas (via a ReplicaSet) and keep that number stable.

Option A is incorrect because `kubectl run` is primarily intended to run a Pod (and in older versions could generate other resources, but it's not the recommended/consistent way to create a Deployment in modern `kubectl` usage). Option C is invalid syntax: `kubectl` subcommand order is incorrect; you don't say `kubectl create nginx deployment`. Option D uses a non-existent `--count` flag for Deployment replicas.

From a Kubernetes fundamentals perspective, this question tests two ideas: (1) Deployments are the standard controller for running stateless workloads with a desired number of replicas, and (2) `kubectl create deployment` is a common imperative shortcut for generating that resource. After running the command, you can confirm with `kubectl get deploy nginx`, `kubectl get rs`, and `kubectl get pods -l app=nginx` (label may vary depending on `kubectl` version). You'll see a ReplicaSet created and two Pods brought up.

In production, teams typically use declarative manifests (`kubectl apply -f`) or GitOps, but knowing the imperative command is useful for quick labs and validation. The key is that replicas are managed by the controller, not by manually starting containers—Kubernetes reconciles the state continuously.

Therefore, B is the verified correct command.

Question: 127

What does "Continuous Integration" mean?

- A. The continuous integration and testing of code changes from multiple sources manually.
- B. The continuous integration and testing of code changes from multiple sources via automation.
- C. The continuous integration of changes from one environment to another.
- D. The continuous integration of new tools to support developers in a project.

Answer: B

Explanation:

The correct answer is B: Continuous Integration (CI) is the practice of frequently integrating code changes from multiple contributors and validating them through automated builds and tests. The “continuous” part is about doing this often (ideally many times per day) and consistently, so integration problems are detected early instead of piling up until a painful merge or release window.

Automation is essential. CI typically includes steps like compiling/building artifacts, running unit and integration tests, executing linters, checking formatting, scanning dependencies for vulnerabilities, and producing build reports. This automation creates fast feedback loops that help developers catch regressions quickly and maintain a releasable main branch.

Option A is incorrect because manual integration/testing does not scale and undermines the reliability and speed that CI is meant to provide. Option C confuses CI with deployment promotion across environments (which is more aligned with Continuous Delivery/Deployment). Option D is unrelated: adding tools can support CI, but it isn't the definition.

In cloud-native application delivery, CI is tightly coupled with containerization and Kubernetes: CI pipelines often build container images from source, run tests, scan images, sign artifacts, and push to registries. Those validated artifacts then flow into CD processes that deploy to Kubernetes using manifests, Helm, or GitOps controllers. Without CI, Kubernetes rollouts become riskier because you lack consistent validation of what you're deploying.

So, CI is best defined as automated integration and testing of code changes from multiple sources, which matches option B.

Question: 128

Which of the following options is true about considerations for large Kubernetes clusters?

- A. Kubernetes supports up to 1000 nodes and recommends no more than 1000 containers per node.
- B. Kubernetes supports up to 5000 nodes and recommends no more than 500 Pods per node.
- C. Kubernetes supports up to 5000 nodes and recommends no more than 110 Pods per node.
- D. Kubernetes supports up to 50 nodes and recommends no more than 1000 containers per node.

Answer: C

Explanation:

The correct answer is C: Kubernetes scalability guidance commonly cites support up to 5000 nodes and recommends no more than 110 Pods per node. The “110 Pods per node” recommendation is a practical limit based on kubelet, networking, and IP addressing constraints, as well as performance characteristics for scheduling, service routing, and node-level resource management. It is also historically aligned with common CNI/IPAM defaults where node Pod CIDRs are sized for ~110 usable Pod IPs.

Why the other options are incorrect: A and D reference “containers per node,” which is not the standard sizing guidance (Kubernetes typically discusses Pods per node). B’s “500 Pods per node” is far above typical recommended limits for many environments and would stress IPAM, kubelet, and node resources significantly.

In large clusters, several considerations matter beyond the headline limits: API server and etcd performance, watch/list traffic, controller reconciliation load, CoreDNS scaling, and metrics/observability overhead. You must also plan for IP addressing (cluster CIDR sizing), node sizes (CPU/memory), and autoscaling behavior. On each node, kubelet and the container runtime must handle churn (starts/stops), logging, and volume operations. Networking implementations (kube-proxy, eBPF dataplanes) also have scaling characteristics.

Kubernetes provides patterns to keep systems stable at scale: request/limit discipline, Pod disruption budgets, topology spread constraints, namespaces and quotas, and careful observability sampling. But the exam-style fact this question targets is the published scalability figure and per-node Pod recommendation.

Therefore, the verified true statement among the options is C.

Question: 129

Which component of the node is responsible to run workloads?

- A. The kubelet.
- B. The kube-proxy.
- C. The kube-apiserver.
- D. The container runtime.

Answer: D

Explanation:

The verified correct answer is D (the container runtime). On a Kubernetes node, the container runtime (such as containerd or CRI-O) is the component that actually executes containers—it creates container processes, manages their lifecycle, pulls images, and interacts with the underlying OS primitives (namespaces, cgroups) through an OCI runtime like runc. In that direct sense, the runtime is what “runs workloads.”

It's important to distinguish responsibilities. The kubelet (A) is the node agent that orchestrates what should run on the node: it watches the API server for Pods assigned to the node and then asks the runtime to start/stop containers accordingly. Kubelet is essential for node management, but it does not itself execute containers; it delegates execution to the runtime via CRI. kube-proxy (B) handles Service traffic routing rules (or is replaced by other dataplanes) and does not run containers. kube-apiserver (C) is a control plane component that stores and serves cluster state; it is not a node workload runner.

So, in the execution chain: scheduler assigns Pod → kubelet sees Pod assigned → kubelet calls runtime via CRI → runtime launches containers. When troubleshooting “containers won't start,” you often inspect kubelet logs and runtime logs because the runtime is the component that can fail image pulls, sandbox creation, or container start operations.

Therefore, the best answer to “which node component is responsible to run workloads” is the container runtime, option D.

Question: 130

The IPv4/IPv6 dual stack in Kubernetes:

- A. Translates an IPv4 request from a Service to an IPv6 Service.
- B. Allows you to access the IPv4 address by using the IPv6 address.
- C. Requires NetworkPolicies to prevent Services from mixing requests.
- D. Allows you to create IPv4 and IPv6 dual stack Services.

Answer: D

Explanation:

The correct answer is D: Kubernetes dual-stack support allows you to create Services (and Pods, depending on configuration) that use both IPv4 and IPv6 addressing. Dual-stack means the cluster is configured to allocate and route traffic for both IP families. For Services, this can mean assigning both an IPv4 ClusterIP and an IPv6 ClusterIP so clients can connect using either family, depending on their network stack and DNS resolution.

Option A is incorrect because dual-stack is not about protocol translation (that would be NAT64/other gateway mechanisms, not the core Kubernetes dual-stack feature). Option B is also a form of translation/aliasing that isn't what Kubernetes dual-stack implies; having both addresses available is different from "access IPv4 via IPv6." Option C is incorrect: dual-stack does not inherently require NetworkPolicies to "prevent mixing requests." NetworkPolicies are about traffic control, not IP family separation.

In Kubernetes, dual-stack requires support across components: the network plugin (CNI) must support IPv4/IPv6, the cluster must be configured with both Pod CIDRs and Service CIDRs, and DNS should return appropriate A and AAAA records for Service names. Once configured, you can specify preferences such as ipFamilyPolicy (e.g., PreferDualStack) and ipFamilies (IPv4, IPv6 order) for Services to influence allocation behavior.

Operationally, dual-stack is useful for environments transitioning to IPv6, supporting IPv6-only clients, or running in mixed networks. But it adds complexity: address planning, firewalling, and troubleshooting need to consider two IP families. Still, the definition in the question is straightforward: Kubernetes dual-stack enables dual-stack Services, which is option D.

Question: 131

What does "continuous" mean in the context of CI/CD?

- A. Frequent releases, manual processes, repeatable, fast processing
- B. Periodic releases, manual processes, repeatable, automated processing
- C. Frequent releases, automated processes, repeatable, fast processing
- D. Periodic releases, automated processes, repeatable, automated processing

Answer: C

Explanation:

The correct answer is C: in CI/CD, “continuous” implies frequent releases, automation, repeatability, and fast feedback/processing. The intent is to reduce batch size and latency between code change and validation/deployment. Instead of integrating or releasing in large, risky chunks, teams integrate changes continually and rely on automation to validate and deliver them safely.

“Continuous” does not mean “periodic” (which eliminates B and D). It also does not mean “manual processes” (which eliminates A and B). Automation is core: build, test, security checks, and deployment steps are consistently executed by pipeline systems, producing reliable outcomes and auditability.

In practice, CI means every merge triggers automated builds and tests so the main branch stays in a healthy state. CD means those validated artifacts are promoted through environments with minimal manual steps, often including progressive delivery controls (canary, blue/green), automated rollbacks on health signal failures, and policy checks. Kubernetes works well with CI/CD because it is declarative and supports rollout primitives: Deployments, readiness probes, and rollback revision history enable safer continuous delivery when paired with pipeline automation.

Repeatability is a major part of “continuous.” The same pipeline should run the same way every time, producing consistent artifacts and deployments. This reduces “works on my machine” issues and shortens incident resolution because changes are traceable and reproducible. Fast processing

and frequent releases also mean smaller diffs, easier debugging, and quicker customer value delivery.

So, the combination that accurately reflects “continuous” in CI/CD is frequent + automated + repeatable + fast, which is option C.

Question: 132

What is ephemeral storage?

- A. Storage space that need not persist across restarts.
- B. Storage that may grow dynamically.
- C. Storage used by multiple consumers (e.g., multiple Pods).

D. Storage that is always provisioned locally.

Answer: A

Explanation:

The correct answer is A: ephemeral storage is non-persistent storage whose data does not need to survive Pod restarts or rescheduling. In Kubernetes, ephemeral storage typically refers to storage tied to the Pod's lifetime—such as the container writable layer, emptyDir volumes, and other temporary storage types. When a Pod is deleted or moved to a different node, that data is generally lost.

This is different from persistent storage, which is backed by PersistentVolumes and PersistentVolumeClaims and is designed to outlive individual Pod instances. Ephemeral storage is commonly used for caches, scratch space, temporary files, and intermediate build artifacts—data that can be recreated and is not the authoritative system of record.

Option B is incorrect because “may grow dynamically” describes an allocation behavior, not the defining characteristic of ephemeral storage. Option C is incorrect because multiple consumers is about access semantics (ReadWriteMany etc.) and shared volumes, not ephemerality. Option D is incorrect because ephemeral storage is not “always provisioned locally” in a strict sense; while many

ephemeral forms are local to the node, the definition is about lifecycle and persistence guarantees, not necessarily physical locality.

Operationally, ephemeral storage is an important scheduling and reliability consideration. Pods can request/limit ephemeral storage similarly to CPU/memory, and nodes can evict Pods under disk pressure. Mismanaged ephemeral storage (logs written to the container filesystem, runaway temp files) can cause node disk exhaustion and cascading failures. Best practices include shipping logs off- node, using emptyDir intentionally with size limits where supported, and using persistent volumes for state that must survive restarts.

So, ephemeral storage is best defined as storage that does not need to persist across restarts/rescheduling, matching option A.

Question: 133

What is the reference implementation of the OCI runtime specification?

- A. lxc
- B. CRI-O
- C. runc
- D. Docker

Answer: C

Explanation:

The verified correct answer is C (runc). The Open Container Initiative (OCI) defines standards for container image format and runtime behavior. The OCI runtime specification describes how to run a container (process execution, namespaces, cgroups, filesystem mounts, capabilities, etc.). runc is widely recognized as the reference implementation of that runtime spec and is used underneath many higher-level container runtimes.

In common container stacks, Kubernetes nodes typically run a CRI-compliant runtime such as

containerd or CRI-O. Those runtimes handle image management, container lifecycle coordination, and CRI integration, but they usually invoke an OCI runtime to actually create and start containers. In many deployments, that OCI runtime is runc (or a compatible alternative). This layering helps keep responsibilities separated: CRI runtime manages orchestration-facing operations; OCI runtime performs the low-level container creation according to the standardized spec.

Option A (lxc) is an older Linux containers technology and tooling ecosystem, but it is not the OCI runtime reference implementation. Option B (CRI-O) is a Kubernetes-focused container runtime that implements CRI; it uses OCI runtimes (often runc) underneath, so it's not the reference implementation itself. Option D (Docker) is a broader platform/tooling suite; while Docker historically used runc under the hood and helped popularize containers, the OCI reference runtime implementation is runc, not Docker.

Understanding this matters in container orchestration contexts because it clarifies what Kubernetes depends on:

Kubernetes relies on CRI for runtime integration, and runtimes rely on OCI standards for interoperability. OCI standards ensure that images and runtime behavior are portable across tools and vendors, and runc is the canonical implementation that demonstrates those standards in practice.

Therefore, the correct answer is C: runc.

Question: 134

What is a Service?

- A. A static network mapping from a Pod to a port.
- B. A way to expose an application running on a set of Pods.
- C. The network configuration for a group of Pods.
- D. An NGINX load balancer that gets deployed for an application.

Answer: B

Explanation:

The correct answer is B: a Kubernetes Service is a stable way to expose an application running on a set of Pods. Pods are ephemeral—IPs can change when Pods are recreated, rescheduled, or scaled. A Service provides a consistent network identity (DNS name and usually a ClusterIP virtual IP) and a policy for routing traffic to the current healthy backends.

Typically, a Service uses a label selector to determine which Pods are part of the backend set. Kubernetes then maintains the corresponding endpoint data (Endpoints/EndpointSlice), and the cluster dataplane (kube-proxy or an eBPF-based implementation) forwards traffic from the Service IP/port to one of the Pod IPs. This enables reliable service discovery and load distribution across replicas, especially during rolling updates where Pods are constantly replaced.

Option A is incorrect because Service routing is not a “static mapping from a Pod to a port.” It’s dynamic and targets a set of Pods. Option C is too vague and misstates the concept; while Services relate to networking, they are not “the network configuration for a group of Pods” (that’s closer to NetworkPolicy/CNI configuration). Option D is incorrect because Kubernetes does not automatically deploy an NGINX load balancer when you create a Service. NGINX might be used as an Ingress controller or external load balancer in some setups, but a Service is a Kubernetes API abstraction, not a specific NGINX component.

Services come in several types (ClusterIP, NodePort, LoadBalancer, ExternalName), but the core definition remains the same: stable access to a dynamic set of Pods. This is foundational for microservices and for decoupling clients from the churn of Pod lifecycles.

So, the verified correct definition is B.

Question: 135

What's the difference between a security profile and a security context?

- A. Security Contexts configure Clusters and Namespaces at runtime. Security profiles are control plane mechanisms to enforce specific settings in the Security Context.
- B. Security Contexts configure Pods and Containers at runtime. Security profiles are control plane mechanisms to enforce specific settings in the Security Context.
- C. Security Profiles configure Pods and Containers at runtime. Security Contexts are control plane mechanisms to enforce specific settings in the Security Profile.
- D. Security Profiles configure Clusters and Namespaces at runtime. Security Contexts are control plane mechanisms to enforce specific settings in the Security Profile.

Answer: B

Explanation:

The correct answer is B. In Kubernetes, a securityContext is part of the Pod and container specification that configures runtime security settings for that workload—things like runAsUser, runAsNonRoot, Linux capabilities, readOnlyRootFilesystem, allowPrivilegeEscalation, SELinux options, seccomp profile selection, and filesystem group (fsGroup). These settings directly affect how the Pod's containers run on the node.

A security profile, in contrast, is a higher-level policy/standard enforced by the cluster control plane (typically via admission control) to ensure workloads meet required security constraints. In modern Kubernetes, this concept aligns with mechanisms like Pod Security Standards (Privileged, Baseline, Restricted) enforced through Pod Security Admission. The "profile" defines what is allowed or forbidden (for example, disallow privileged containers, disallow hostPath mounts, require non-root, restrict capabilities). The control plane enforces these constraints by validating or rejecting Pod specs that do not comply—ensuring consistent security posture across namespaces and teams.

Option A and D are incorrect because security contexts do not "configure clusters and namespaces at runtime"; security contexts apply to Pods/containers. Option C reverses the relationship: security profiles don't configure Pods at runtime; they constrain what security context settings (and other fields) are acceptable.

Practically, you can think of it as:

SecurityContext = workload-level configuration knobs (declared in manifests, applied at runtime).

SecurityProfile/Standards = cluster-level guardrails that determine which knobs/settings are permitted.

This separation supports least privilege: developers declare needed runtime settings, and cluster governance ensures those settings stay within approved boundaries. Therefore, B is the verified answer.

Question: 136

At which layer would distributed tracing be implemented in a cloud native deployment?

- A. Network
- B. Application
- C. Database
- D. Infrastructure

Answer: B

Explanation:

Distributed tracing is implemented primarily at the application layer, so B is correct. The reason is simple: tracing is about capturing the end-to-end path of a request as it traverses services, libraries, queues, and databases. That "request context" (trace ID, span ID, baggage) must be created, propagated, and enriched as code executes. While infrastructure components (proxies, gateways, service meshes) can generate or augment trace spans, the fundamental unit of tracing is still tied to application operations (an HTTP handler, a gRPC call, a database query, a cache lookup).

In Kubernetes-based microservices, distributed tracing typically uses standards like OpenTelemetry for instrumentation and context propagation. Application frameworks emit spans for key operations, attach attributes (route, status code, tenant, retry count), and propagate context via headers (e.g., W3C Trace Context). This is what lets you reconstruct "Service A → Service B → Service C" for one user request and identify the slow or failing hop.

Why other layers are not the best answer:

Network focuses on packets/flows, but tracing is not a packet-capture problem; it's a causal requestpath problem across services.

Database spans are part of traces, but tracing is not "implemented in the database layer" overall; DB spans are one

component.

Infrastructure provides the platform and can observe traffic, but without application context it can't fully represent business operations (and many useful attributes live in app code).

So the correct layer for "where tracing is implemented" is the application layer—even when a mesh or proxy helps, it's still describing application request execution across components.

Question: 137

What framework does Kubernetes use to authenticate users with JSON Web Tokens?

- A. OpenID Connect
- B. OpenID Container
- C. OpenID Cluster
- D. OpenID CNCF

Answer: A

Explanation:

Kubernetes commonly authenticates users using OpenID Connect (OIDC) when JSON Web Tokens (JWTs) are involved, so A is correct. OIDC is an identity layer on top of OAuth 2.0 that standardizes how clients obtain identity information and how JWTs are issued and validated.

In Kubernetes, authentication happens at the API server. When OIDC is configured, the API server validates incoming bearer tokens (JWTs) by checking token signature and claims against the configured OIDC issuer and client settings. Kubernetes can use OIDC claims (such as sub, email, groups) to map the authenticated identity to Kubernetes RBAC subjects. This is how enterprises integrate clusters with identity providers such as Okta, Dex, Azure AD, or other OIDC-compliant IdPs.

Options B, C, and D are fabricated phrases and not real frameworks. Kubernetes documentation explicitly references OIDC as a supported method for token-based user authentication (alongside client certificates, bearer tokens, static token files, and webhook authentication). The key point is that Kubernetes does not "invent" JWT auth; it integrates with standard identity providers through OIDC so clusters can participate in centralized SSO and group-based authorization.

Operationally, OIDC authentication is typically paired with:

RBAC for authorization (“what you can do”)

Audit logging for traceability

Short-lived tokens and rotation practices for security

Group claim mapping to simplify permission management

So, the verified framework Kubernetes uses with JWTs for user authentication is OpenID Connect.

Question: 138

Which of the following is a feature Kubernetes provides by default as a container orchestration tool?

- A. A portable operating system.
- B. File system redundancy.
- C. A container image registry.
- D. Automated rollouts and rollbacks.

Answer: D

Explanation:

Kubernetes provides automated rollouts and rollbacks for workloads by default (via controllers like Deployments), so D is correct. In Kubernetes, application delivery is controller-driven: you declare the desired state (new image, new config), and controllers reconcile the cluster toward that state. Deployments implement rolling updates, gradually replacing old Pods with new ones while respecting availability constraints. Kubernetes tracks rollout history and supports rollback to previous ReplicaSets when an update fails or is deemed unhealthy.

This is a core orchestration capability: it reduces manual intervention and makes change safer. Rollouts use readiness checks and update strategies to avoid taking the service down, and `kubectl rollout status/history/undo` supports day-to-day release operations.

The other options are not “default Kubernetes orchestration features”:

Kubernetes is not a portable operating system (A). It's a platform for orchestrating containers on top of an OS.

Kubernetes does not provide filesystem redundancy by itself (B). Storage redundancy is handled by underlying storage systems and CSI drivers (e.g., replicated block storage, distributed filesystems).

Kubernetes does not include a built-in container image registry (C). You use external registries (Docker Hub, ECR, GCR, Harbor, etc.). Kubernetes pulls images but does not host them as a core feature.

So the correct "provided by default" orchestration feature in this list is the ability to safely manage application updates via automated rollouts and rollbacks.

Question: 139

Which of the following sentences is true about container runtimes in Kubernetes?

- A. If you let iptables see bridged traffic, you don't need a container runtime.
- B. If you enable IPv4 forwarding, you don't need a container runtime.
- C. Container runtimes are deprecated, you must install CRI on each node.
- D. You must install a container runtime on each node to run pods on it.

Answer: D

Explanation:

A Kubernetes node must have a container runtime to run Pods, so D is correct. Kubernetes schedules Pods to nodes, but the actual execution of containers is performed by a runtime such as containerd or CRI-O. The kubelet communicates with that runtime via the Container Runtime Interface (CRI) to pull images, create sandboxes, and start/stop containers. Without a runtime, the node cannot launch container processes, so Pods cannot transition into running state.

Options A and B confuse networking kernel settings with runtime requirements. iptables bridged

traffic visibility and IPv4 forwarding can be relevant for node networking, but they do not replace the need for a container runtime. Networking and container execution are separate layers: you need networking for connectivity, and

you need a runtime for running containers.

Option C is also incorrect and muddled. Container runtimes are not deprecated; rather, Kubernetes removed the built-in Docker shim integration from kubelet in favor of CRI-native runtimes. CRI is an interface, not “something you install instead of a runtime.” In practice you install a CRI-compatible runtime (containerd/CRI-O), which implements CRI endpoints that kubelet talks to.

Operationally, the runtime choice affects node behavior: image management, logging integration, performance characteristics, and compatibility. Kubernetes installation guides explicitly list installing a container runtime as a prerequisite for worker nodes. If a cluster has nodes without a properly configured runtime, workloads scheduled there will fail to start (often stuck in ContainerCreating/ImagePullBackOff/Runtime errors).

Therefore, the only fully correct statement is D: each node needs a container runtime to run Pods.

Question: 140

If a Pod was waiting for container images to download on the scheduled node, what state would it be in?

- A. Failed
- B. Succeeded
- C. Unknown
- D. Pending

Answer: D

Explanation:

If a Pod is waiting for its container images to be pulled to the node, it remains in the Pending phase, so D is correct. Kubernetes Pod “phase” is a high-level summary of where the Pod is in its lifecycle.

Pending means the Pod has been accepted by the cluster but one or more of its containers has not started yet. That can occur because the Pod is waiting to be scheduled, waiting on volume attachment/mount, or—very commonly—waiting for the container runtime to pull the image.

When image pulling is the blocker, kubectl describe pod <name> usually shows events like "Pulling image ..." and "Successfully pulled image ..." or failures like ImagePullBackOff/ErrImagePull. Even if the node has been assigned (scheduler has set spec.nodeName), the Pod can still be Pending while kubelet and the runtime perform preparation steps.

Why the other phases don't apply:

Succeeded is for run-to-completion Pods that have finished successfully (typical for Jobs).

Failed means the Pod terminated and at least one container terminated in failure (and won't be restarted, depending on restartPolicy).

Unknown is used when the node can't be contacted and the Pod's state can't be reliably determined (rare in healthy clusters).

A subtle but important Kubernetes detail: status "Waiting" reasons like ImagePullBackOff are container states inside .status.containerStatuses, while the Pod phase can still be Pending. So, "waiting for images to download" maps to Pod Pending, with container waiting reasons providing the deeper diagnosis.

Therefore, the verified correct answer is D: Pending.

Question: 141

What is CloudEvents?

- A. It is a specification for describing event data in common formats for Kubernetes network traffic management and cloud providers.
- B. It is a specification for describing event data in common formats in all cloud providers including major cloud providers.
- C. It is a specification for describing event data in common formats to provide interoperability across services, platforms and systems.
- D. It is a Kubernetes specification for describing events data in common formats for iCloud services, iOS platforms and iMac.

Answer: C

Explanation:

CloudEvents is an open specification for describing event data in a common way to enable interoperability across services, platforms, and systems, so C is correct. In cloud-native architectures, many components communicate asynchronously via events (message brokers, event buses, webhooks). Without a standard envelope, each producer and consumer invents its own event structure, making integration brittle. CloudEvents addresses this by standardizing core metadata fields—like event id, source, type, specversion, and time—and defining how event payloads are carried.

This helps systems interoperate regardless of transport. CloudEvents can be serialized as JSON or other encodings and carried over HTTP, messaging systems, or other protocols. By using a shared spec, you can route, filter, validate, and transform events more consistently.

Option A is too narrow and incorrectly ties CloudEvents to Kubernetes traffic management; CloudEvents is broader than Kubernetes. Option B is closer but still framed incorrectly—CloudEvents is not merely “for all cloud providers,” it is an interoperability spec across services and platforms, including but not limited to cloud provider event systems. Option D is clearly incorrect.

In Kubernetes ecosystems, CloudEvents is relevant to event-driven systems and serverless platforms (e.g., Knative Eventing and other eventing frameworks) because it provides a consistent event contract across producers and consumers. That consistency reduces coupling, supports better tooling (schema validation, tracing correlation), and makes event-driven architectures easier to operate at scale.

So, the correct definition is C: a specification for common event formats to enable interoperability across systems.

Question: 142

What is the main purpose of etcd in Kubernetes?

- A. etcd stores all cluster data in a key value store.
- B. etcd stores the containers running in the cluster for disaster recovery.
- C. etcd stores copies of the Kubernetes config files that live /etc/.
- D. etcd stores the YAML definitions for all the cluster components.

Answer: A

Explanation:

The main purpose of etcd in Kubernetes is to store the cluster's state as a distributed key-value store, so A is correct. Kubernetes is API-driven: objects like Pods, Deployments, Services, ConfigMaps, Secrets, Nodes, and RBAC rules are persisted by the API server into etcd. Controllers, schedulers, and other components then watch the API for changes and reconcile the cluster accordingly. This makes etcd the "source of truth" for desired and observed cluster state.

Options B, C, and D are misconceptions. etcd does not store the running containers; that's the job of the kubelet/container runtime on each node, and container state is ephemeral. etcd does not store /etc configuration file copies. And while you may author objects as YAML manifests, Kubernetes stores them internally as API objects (serialized) in etcd—not as "YAML definitions for all components." The data is structured key/value entries representing Kubernetes resources and metadata.

Because etcd is so critical, its performance and reliability directly affect the cluster. Slow disk I/O or poor network latency increases API request latency and can delay controller reconciliation, leading to cascading operational problems (slow rollouts, delayed scheduling, timeouts). That's why etcd is typically run on fast, reliable storage and in an HA configuration (often 3 or 5 members) to maintain quorum and tolerate failures. Backups (snapshots) and restore procedures are also central to disaster recovery: if etcd is lost, the cluster loses its state.

Security is also important: etcd can contain sensitive information (especially Secrets unless encrypted at rest). Proper TLS, restricted access, and encryption-at-rest configuration are standard best practices.

So, the verified correct answer is A: etcd stores all cluster data/state in a key-value store.

Question: 143

Imagine you're releasing open-source software for the first time. Which of the following is a valid semantic version?

- A. 1.0
- B. 2021-10-11
- C. 0.1.0-rc
- D. v1beta1

Answer: C

Explanation:

Semantic Versioning (SemVer) follows the pattern MAJOR.MINOR.PATCH with optional pre-release identifiers (e.g., -rc, -alpha.1) and build metadata. Among the options, 0.1.0-rc matches SemVer rules, so C is correct.

0.1.0-rc breaks down as: MAJOR=0, MINOR=1, PATCH=0, and -rc indicates a pre-release ("release candidate"). Pre-release versions are valid SemVer and are explicitly allowed to denote versions that are not yet considered stable. For a first-time open-source release, 0.x.y is common because it signals the API may still change in backward-incompatible ways before reaching 1.0.0.

Why the other options are not correct SemVer as written:

1 .0 is missing the PATCH segment; SemVer requires three numeric components (e.g., 1.0.0).

2021-10-11 is a date string, not MAJOR.MINOR.PATCH.

v1beta1 resembles Kubernetes API versioning conventions, not SemVer.

In cloud-native delivery and Kubernetes ecosystems, SemVer matters because it communicates compatibility. Incrementing MAJOR indicates breaking changes, MINOR indicates backward-compatible feature additions, and PATCH indicates backward-compatible bug fixes. Pre-release tags allow releasing candidates for testing without claiming full stability. This is especially useful for opensource consumers and automation systems that need consistent version comparison and upgrade planning.

50 , the only valid semantic version in the choices is 0.1.0-rc, option C.

Question: 144

Which cloud native tool keeps Kubernetes clusters in sync with sources of configuration (like Git repositories), and automates updates to configuration when there is new code to deploy?

- A. Flux and ArgoCD
- B. GitOps Toolkit
- C. Linkerd and Istio
- D. Helm and Kustomize

Answer: A

Explanation:

Tools that continuously reconcile cluster state to match a Git repository's desired configuration are GitOps controllers,

and the best match here is Flux and ArgoCD, so A is correct. GitOps is the practice where Git is the source of truth for declarative system configuration. A GitOps tool continuously compares the desired state (manifests/Helm/Kustomize outputs stored in Git) with the actual state in the cluster and then applies changes to eliminate drift.

Flux and Argo CD both implement this reconciliation loop. They watch Git repositories, detect updates (new commits/tags), and apply the updated Kubernetes resources. They also surface drift and sync status, enabling auditable, repeatable deployments and easy rollbacks (revert Git). This model improves delivery velocity and security because changes flow through code review, and cluster changes can be restricted to the GitOps controller identity rather than ad-hoc human kubectl access.

Option B ("GitOps Toolkit") is related—Flux uses a GitOps Toolkit internally—but the question asks for a "tool" that keeps clusters in sync; the recognized tools are Flux and Argo CD in this list. Option C lists service meshes (traffic/security/telemetry), not deployment synchronization tools. Option D lists packaging/templating tools; Helm and Kustomize help build manifests, but they do not, by themselves, continuously reconcile cluster state to a Git source.

In Kubernetes application delivery, GitOps tools become the deployment engine: CI builds artifacts, updates references in Git (image tags/digests), and the GitOps controller deploys those changes. This separation strengthens traceability and reduces configuration drift. Therefore, A is the verified correct answer.

Question: 145

Which tool is used to streamline installing and managing Kubernetes applications?

- A. apt
- B. helm
- C. service
- D. brew

Answer: B

Explanation:

Helm is the Kubernetes package manager used to streamline installing and managing applications, so B is correct. Helm packages Kubernetes resources into charts, which contain templates, default values, and metadata. When you install a chart, Helm renders templates into concrete manifests and applies them to the cluster. Helm also tracks a "release," enabling upgrades, rollbacks, and consistent lifecycle operations across environments.

This is why Helm is widely used for complex applications that require multiple Kubernetes objects (Deployments/StatefulSets, Services, Ingresses, ConfigMaps, RBAC, CRDs). Rather than manually maintaining many YAML files per environment, teams can parameterize configuration with values and reuse the same chart across dev/stage/prod with different overrides.

Option A (apt) and option D (brew) are OS package managers (Debian/Ubuntu and macOS/Linuxbrew respectively), not Kubernetes application managers. Option C (service) is a Linux service manager command pattern and not relevant here.

In cloud-native delivery pipelines, Helm often integrates with GitOps and CI/CD: the pipeline builds an image, updates chart values (image tag/digest), and deploys via Helm or via GitOps controllers that render/apply Helm charts. Helm also supports chart repositories and versioning, making it easier to standardize deployments and manage dependencies.

So, the verified tool for streamlined Kubernetes app install/management is Helm (B).

Question: 146

What are the two steps performed by the kube-scheduler to select a node to schedule a pod?

- A. Grouping and placing
- B. Filtering and selecting
- C. Filtering and scoring
- D. Scoring and creating

Answer: C

Explanation:

The kube-scheduler selects a node in two main phases: filtering and scoring, so C is correct. First, filtering identifies which nodes are feasible for the Pod by applying hard constraints. These include resource availability (CPU/memory requests), node taints/tolerations, node selectors and required affinities, topology constraints, and other scheduling requirements.

Nodes that cannot satisfy the Pod's requirements are removed from consideration.

Second, scoring ranks the remaining feasible nodes using priority functions to choose the "best" placement. Scoring can consider factors like spreading Pods across nodes/zones, packing efficiency, affinity preferences, and other policies configured in the scheduler. The node with the highest score is selected (with tie-breaking), and the scheduler binds the Pod by setting spec.nodeName.

Option B ("filtering and selecting") is close but misses the explicit scoring step that is central to scheduler design. The scheduler does "select" a node, but the canonical two-step wording in Kubernetes scheduling is filtering then scoring.

Options A and D are not how scheduler internals are

described.

Operationally, understanding filtering vs scoring helps troubleshoot scheduling failures. If a Pod can't be scheduled, it failed in filtering—`kubectl describe pod` often shows "0/... nodes are available" reasons (insufficient CPU, taints, affinity mismatch). If it schedules but lands in unexpected places, it's often about scoring preferences (affinity weights, topology spread preferences, default scheduler profiles).

So the verified correct answer is C: kube-scheduler uses Filtering and Scoring.

Question: 147

To visualize data from Prometheus you can use expression browser or console templates. What is the other data visualization tool commonly used together with Prometheus?

- A. Grafana
- B. Graphite
- C. Nirvana
- D. GraphQL

Answer: A

Explanation:

The most common visualization tool used with Prometheus is Grafana, so A is correct. Prometheus includes a built-in expression browser that can graph query results, but Grafana provides a much richer dashboarding experience: reusable dashboards, variables, templating, annotations, alerting integrations, and multi-data-source support.

In Kubernetes observability stacks, Prometheus scrapes and stores time-series metrics (cluster and application metrics). Grafana queries Prometheus using PromQL and renders the results into dashboards for SREs and developers. This pairing is widespread because it cleanly separates concerns: Prometheus is the metrics store and query engine; Grafana is the UI and dashboard layer.

Option B (Graphite) is a separate metrics system with its own storage/query model; while Grafana can visualize Graphite too, the question asks what is commonly used together with Prometheus, which is Grafana. Option D (GraphQL) is an API query language, not a metrics visualization tool. Option C ("Nirvana") is not a standard Prometheus visualization tool in common Kubernetes stacks.

In practice, this combo enables operational outcomes: dashboards for error rates and latency (often derived from histograms), capacity monitoring (node CPU/memory), workload behavior (Pod restarts, HPA scaling), and SLO reporting. Grafana dashboards often serve as the shared language during incidents: teams correlate alerts with time-series patterns and quickly identify when regressions began.

Therefore, the verified correct tool commonly used with Prometheus for visualization is Grafana (A).

Question: 148

Which item is a Kubernetes node component?

- A. kube-scheduler
- B. kubectI
- C. kube-proxy
- D. etcd

Answer: C

Explanation:

A Kubernetes node component is a component that runs on worker nodes to support Pods and node-level networking/operations. Among the options, kube-proxy is a node component, so C is correct.

kube-proxy runs on each node and implements parts of the Kubernetes Service networking model. It watches the API server for Service and endpoint updates and then programs node networking rules (iptables/IPVS, or equivalent) so traffic sent to a Service IP/port is forwarded to one of the backend Pod endpoints. This is essential for stable virtual IPs and load distribution across Pods.

Why the other options are not node components:

kube-scheduler is a control plane component; it assigns Pods to nodes but does not run on every node as part of node functionality.

kubectl is a client CLI tool used by humans/automation; it is not a cluster component.

etcd is the control plane datastore; it stores cluster state and is not a per-node workload component.

Operationally, kube-proxy can be replaced by some modern CNI/eBPF dataplanes, but in classic Kubernetes architecture it remains the canonical node-level component for Service rule programming. Understanding which components are node vs control plane is key for troubleshooting: node issues involve kubelet/runtime/kube-proxy/CNI; control plane issues involve API server/scheduler/controller-manager/etcd.

So, the verified node component in this list is kube-proxy (C).

Question: 149

In a serverless computing architecture:

- A. Users of the cloud provider are charged based on the number of requests to a function.
- B. Serverless functions are incompatible with containerized functions.
- C. Users should make a reservation to the cloud provider based on an estimation of usage.
- D. Containers serving requests are running in the background in idle status.

Answer: A

Explanation:

Serverless architectures typically bill based on actual consumption, often measured as number of requests and execution duration (and sometimes memory/CPU allocated), so A is correct. The defining trait is that you don't provision or manage servers directly; the platform scales execution up and down automatically, including down to zero for many models, and charges you for what you use.

Option B is incorrect: many serverless platforms can run container-based workloads (and some are explicitly "serverless containers"). The idea is the operational abstraction and billing model, not incompatibility with containers. Option C is incorrect because "making a reservation based on estimation" describes reserved capacity purchasing, which is the opposite of the typical serverless pay-per-use model. Option D is misleading: serverless systems aim to avoid charging for idle compute; while platforms may keep some warm capacity for latency reasons, the customer-facing model is not "containers running idle in the background."

In cloud-native architecture, serverless is often chosen for spiky, event-driven workloads where you want minimal ops overhead and cost efficiency at low utilization. It pairs naturally with eventing systems (queues, pub/sub) and can be integrated with Kubernetes ecosystems via event-driven autoscaling frameworks or managed serverless offerings.

So the correct statement is A: charging is commonly based on requests (and usage), which captures the cost and operational model that differentiates serverless from always-on infrastructure.

Question: 150

How is application data maintained in containers?

- A. Store data into data folders.
- B. Store data in separate folders.
- C. Store data into sidecar containers.
- D. Store data into volumes.

Answer: D

Explanation:

Container filesystems are ephemeral: the writable layer is tied to the container lifecycle and can be lost when containers are recreated. Therefore, maintaining application data correctly means storing it in volumes, making D the correct answer. In Kubernetes, volumes provide durable or shareable storage that is mounted into containers at specific paths. Depending on the volume type, the data can persist across container restarts and even Pod rescheduling.

Kubernetes supports many volume patterns. For transient scratch data you might use emptyDir (ephemeral for the Pod's lifetime). For durable state, you typically use PersistentVolumes consumed by PersistentVolumeClaims (PVCs), backed by storage systems via CSI drivers (cloud disks, SAN/NAS, distributed storage). This decouples the application container image from its state and enables rolling updates, rescheduling, and scaling without losing data.

Options A and B ("folders") are incomplete because folders inside the container filesystem do not guarantee persistence. A folder is only as durable as the underlying storage; without a mounted volume, it lives in the container's writable layer and will disappear when the container is replaced. Option C is incorrect because "sidecar containers" are not a data durability mechanism; sidecars can help ship logs or sync data, but persistent data should still be stored on volumes (or external services like managed databases).

From an application delivery standpoint, the principle is: containers should be immutable and disposable, and state should be externalized. Volumes (and external managed services) make this possible. In Kubernetes, this is a foundational pattern enabling safe rollouts, self-healing, and portability: the platform can kill and recreate Pods freely because data is maintained independently via volumes.

Therefore, the verified correct choice is D: Store data into volumes.

Question: 151

Which of the following scenarios would benefit the most from a service mesh architecture?

- A. A few applications with hundreds of Pod replicas running in multiple clusters, each one providing multiple services.
- B. Thousands of distributed applications running in a single cluster, each one providing multiple services.
- C. Tens of distributed applications running in multiple clusters, each one providing multiple services.

D. Thousands of distributed applications running in multiple clusters, each one providing multiple services.

Answer: D

Explanation:

A service mesh is most valuable when service-to-service communication becomes complex at large scale—many services, many teams, and often multiple clusters. That’s why D is the best fit: thousands of distributed applications across multiple clusters. In that scenario, the operational burden of securing, observing, and controlling east-west traffic grows dramatically. A service mesh (e.g., Istio, Linkerd) addresses this by introducing a dedicated networking layer (usually sidecar proxies such as Envoy) that standardizes capabilities across services without requiring each application to implement them consistently.

The common “mesh” value-adds are: mTLS for service identity and encryption, fine-grained traffic policy (retries, timeouts, circuit breaking), traffic shifting (canary, mirroring), and consistent telemetry (metrics, traces, access logs). Those features become increasingly beneficial as the number of services and cross-service calls rises, and as you add multi-cluster routing, failover, and policy management across environments. With thousands of applications, inconsistent libraries and configurations become a reliability and security risk; the mesh centralizes and standardizes these behaviors.

In smaller environments (A or C), you can often meet requirements with simpler approaches: Kubernetes Services, Ingress/Gateway, basic mTLS at the edge, and application-level libraries. A single large cluster (B) can still benefit from a mesh, but adding multiple clusters increases complexity: traffic management across clusters, identity trust domains, global observability correlation, and consistent policy enforcement. That’s where mesh architectures typically justify their additional overhead (extra proxies, control plane components, operational complexity).

So, the “most benefit” scenario is the largest, most distributed footprint—D.

Question: 152

Kubernetes protect you against voluntary interruptions (such as deleting Pods, draining nodes) to run applications in a highly available manner.

A. Pod Topology Spread Constraints

B. Pod Disruption Budgets

C. Taints and Tolerations

D. Resource Limits and Requests

Answer: B

Explanation:

The correct answer is B: Pod Disruption Budgets (PDBs). A PDB is a policy object that limits how many Pods of an application can be voluntarily disrupted at the same time. "Voluntary disruptions" include actions such as draining a node for maintenance (kubectl drain), cluster upgrades, or an administrator deleting Pods. The core purpose is to preserve availability by ensuring that a minimum number (or percentage) of replicas remain running and ready while those planned disruptions occur.

A PDB is typically defined with either minAvailable (e.g., "at least 3 Pods must remain available") or maxUnavailable (e.g., "no more than 1 Pod can be unavailable"). Kubernetes uses this budget when performing eviction operations. If evicting a Pod would violate the PDB, the eviction is blocked (or delayed), which forces maintenance workflows to proceed more safely—either by draining more slowly, scaling up first, or scheduling maintenance in stages.

Why the other options are not correct: topology spread constraints (A) influence scheduling distribution across failure domains but don't directly protect against voluntary disruptions. Taints and tolerations (C) control where Pods can schedule, not how many can be disrupted. Resource requests/limits (D) control CPU/memory allocation and do not guard availability during drains or deletions.

PDBs also work best when paired with Deployments/StatefulSets that maintain replicas and with readiness probes that accurately represent whether a Pod can serve traffic. PDBs do not prevent involuntary disruptions (node crashes), but they materially reduce risk during planned operations— exactly what the question is targeting.

Question: 153

What sentence is true about CronJobs in Kubernetes?

- A. A CronJob creates one or multiple Jobs on a repeating schedule.
- B. A CronJob creates one container on a repeating schedule.
- C. CronJobs are useful on Linux but are obsolete in Kubernetes.
- D. The CronJob schedule format is different in Kubernetes and Linux.

Answer: A

Explanation:

The true statement is A: a Kubernetes CronJob creates Jobs on a repeating schedule. CronJob is a controller designed for time-based execution. You define a schedule using standard cron syntax (minute, hour, day-of-month, month, day-of-week), and when the schedule triggers, the CronJob controller creates a Job object. Then the Job controller creates one or more Pods to run the task to completion.

Option B is incorrect because CronJobs do not “create one container”; they create Jobs, and Jobs create Pods (which may contain one or multiple containers). Option C is wrong because CronJobs are a core Kubernetes workload primitive for recurring tasks and remain widely used for periodic work like backups, batch processing, and cleanup. Option D is wrong because Kubernetes CronJobs intentionally use cron-like scheduling expressions; the format aligns with the cron concept (with Kubernetes-specific controller behavior around missed runs, concurrency, and history).

CronJobs also provide operational controls you don't get from plain Linux cron on a node:

concurrencyPolicy (Allow/Forbid/Replace) to manage overlapping runs

startingDeadlineSeconds to control how missed schedules are handled

history limits for successful/failed Jobs to avoid clutter

integration with Kubernetes RBAC, Secrets, ConfigMaps, and volumes for consistent runtime configuration

consistent execution environment via container images, not ad-hoc node scripts

Because the CronJob creates Jobs as first-class API objects, you get observability (events/status), predictable retries, and lifecycle management. That's why the accurate statement is A.

Question: 154

What is the purpose of the kube-proxy?

- A. The kube-proxy balances network requests to Pods.
- B. The kube-proxy maintains network rules on nodes.
- C. The kube-proxy ensures the cluster connectivity with the internet.
- D. The kube-proxy maintains the DNS rules of the cluster.

Answer: B

Explanation:

The correct answer is B: kube-proxy maintains network rules on nodes. kube-proxy is a node component that implements part of the Kubernetes Service abstraction. It watches the Kubernetes API for Service and EndpointSlice/Endpoints changes, and then programs the node's dataplane rules (commonly iptables or IPVS, depending on configuration) so that traffic sent to a Service virtual IP and port is correctly forwarded to one of the backing Pod endpoints.

This is how Kubernetes provides stable Service addresses even though Pod IPs are ephemeral. When Pods scale up/down or are replaced during a rollout, endpoints change; kube-proxy updates the node rules accordingly. From the perspective of a client, the Service name and ClusterIP remain stable, while the actual backend endpoints are load-distributed.

Option A is a tempting phrasing but incomplete: load distribution is an outcome of the forwarding rules, but kube-proxy's primary role is maintaining the network forwarding rules that make Services work. Option C is incorrect because internet connectivity depends on cluster networking, routing, NAT, and often CNI configuration—not kube-proxy's job description. Option D is incorrect because DNS is typically handled by CoreDNS; kube-proxy does not "maintain DNS rules."

Operationally, kube-proxy failures often manifest as Service connectivity issues: Pod-to-Service traffic fails, ClusterIP routing breaks, NodePort behavior becomes inconsistent, or endpoints aren't updated correctly. Modern Kubernetes environments sometimes replace kube-proxy with eBPF-based dataplanes, but in the classic architecture the correct statement remains: kube-proxy runs on each node and maintains the rules needed for Service traffic steering.

Question: 155

Manual reclamation policy of a PV resource is known as:

- A. claimRef
- B. Delete
- C. Retain
- D. Recycle

Answer: C

Explanation:

The correct answer is C: Retain. In Kubernetes persistent storage, a PersistentVolume (PV) has a persistentVolumeReclaimPolicy that determines what happens to the underlying storage asset after its PersistentVolumeClaim (PVC) is deleted. The reclaim policy options historically include Delete and Retain (and Recycle, which is deprecated/removed in many modern contexts). “Manual reclamation” refers to the administrator having to manually clean up and/or rebind the storage after the claim is released—this behavior corresponds to Retain.

With Retain, when the PVC is deleted, the PV moves to a “Released” state, but the actual storage resource (cloud disk, NFS path, etc.) is not deleted automatically. Kubernetes will not automatically make that PV available for a new claim until an administrator takes action—typically cleaning the data, removing the old claim reference, and/or creating a new PV/PVC binding flow. This is important for data safety: you don’t want to automatically delete sensitive or valuable data just because a claim was removed.

By contrast, Delete means Kubernetes (via the storage provisioner/CSI driver) will delete the underlying storage asset when the claim is deleted—useful for dynamic provisioning and disposable environments. Recycle used to scrub the volume contents and make it available again, but it’s not the recommended modern approach and has been phased out in favor of dynamic provisioning and explicit workflows.

So, the policy that implies manual intervention and manual cleanup/reuse is Retain, which is option C.

Question: 156

Which component of the Kubernetes architecture is responsible for integration with the CRI container runtime?

- A. kubeadm
- B. kubelet
- C. kube-apiserver
- D. kubect1

Answer: B

Explanation:

The correct answer is B: kubelet. The Container Runtime Interface (CRI) defines how Kubernetes interacts with container runtimes in a consistent, pluggable way. The component that speaks CRI is the kubelet, the node agent responsible for running Pods on each node. When the kube-scheduler assigns a Pod to a node, the kubelet reads the PodSpec and makes the runtime calls needed to realize that desired state—pull images, create a Pod sandbox, start containers, stop containers, and retrieve status and logs. Those calls are made via CRI to a CRI-compliant runtime such as containerd or CRI-O.

Why not the others:

kubeadm bootstraps clusters (init/join/upgrade workflows) but does not run containers or speak CRI for workload execution.

kube-apiserver is the control plane API frontend; it stores and serves cluster state and does not directly integrate with runtimes.

kubectrl is just a client tool that sends API requests; it is not involved in runtime integration on nodes.

This distinction matters operationally. If the runtime is misconfigured or CRI endpoints are unreachable, kubelet will report errors and Pods can get stuck in ContainerCreating, image pull failures, or runtime errors. Debugging often involves checking kubelet logs and runtime service health, because kubelet is the integration point bridging Kubernetes scheduling/state with actual container execution.

So, the node-level component responsible for CRI integration is the kubelet—option B.

Question: 157

Which one of the following is an open source runtime security tool?

- A. lxd
- B. containerd
- C. falco

D. gVisor

Answer: C

Explanation:

The correct answer is C: Falco. Falco is a widely used open-source runtime security tool (originally created by Sysdig and now a CNCF project) designed to detect suspicious behavior at runtime by monitoring system calls and other kernel-level signals. In Kubernetes environments, Falco helps identify threats such as unexpected shell access in containers, privilege escalation attempts, access to sensitive files, anomalous network tooling, crypto-mining patterns, and other behaviors that indicate compromise or policy violations.

The other options are not primarily “runtime security tools” in the detection/alerting sense: containerd is a container runtime responsible for executing containers; it’s not a security detection tool.

lxd is a system container and VM manager; again, not a runtime threat detection tool.

gVisor is a sandboxed container runtime that improves isolation by interposing a user-space kernel; it’s a security mechanism, but the question asks for a runtime security tool (monitoring/detection). Falco fits that definition best.

In cloud-native security practice, Falco typically runs as a DaemonSet so it can observe activity on every node. It uses rules to define what “bad” looks like and can emit alerts to SIEM systems, logging backends, or incident response workflows. This complements preventative controls like RBAC, Pod Security Admission, seccomp, and least privilege configurations. Preventative controls reduce risk; Falco provides visibility and detection when something slips through.

Therefore, among the provided choices, the verified runtime security tool is Falco (C).

Question: 158

What are the advantages of adopting a GitOps approach for your deployments?

A. Reduce failed deployments, operational costs, and fragile release processes.

- B. Reduce failed deployments, configuration drift, and fragile release processes.
- C. Reduce failed deployments, operational costs, and learn git.
- D. Reduce failed deployments, configuration drift and improve your reputation.

Answer: B

Explanation:

The correct answer is B: GitOps helps reduce failed deployments, reduce configuration drift, and reduce fragile release processes. GitOps is an operating model where Git is the source of truth for declarative configuration (Kubernetes manifests, Helm releases, Kustomize overlays). A GitOps controller (like Flux or Argo CD) continuously reconciles the cluster's actual state to match what's

declared in Git. This creates a stable, repeatable deployment pipeline and minimizes "snowflake" environments.

Reducing failed deployments: changes go through pull requests, code review, automated checks, and controlled merges.

Deployments become predictable because the controller applies known-good, versioned configuration rather than ad-hoc manual commands. Rollbacks are also simpler—reverting a Git commit returns the cluster to the prior desired state.

Reducing configuration drift: without GitOps, clusters often drift because humans apply hotfixes directly in production or because different environments diverge over time. With GitOps, the controller detects drift and either alerts or automatically corrects it, restoring alignment with Git.

Reducing fragile release processes: releases become standardized and auditable. Git history provides an immutable record of who changed what and when. Promotion between environments becomes systematic (merge/branch/tag), and the same declarative artifacts are used consistently.

The other options include items that are either not the primary GitOps promise (like "learn git") or subjective ("improve your reputation"). Operational cost reduction can happen indirectly through fewer incidents and more automation, but the most canonical and direct GitOps advantages in Kubernetes delivery are reliability and drift control—captured precisely in B.

Question: 159

Which command lists the running containers in the current Kubernetes namespace?

- A. kubectl get pods
- B. kubectl ls
- C. kubectl ps
- D. kubectl show pods

Answer: A

Explanation:

The correct answer is A: kubectl get pods. Kubernetes does not manage “containers” as standalone top-level objects; the primary schedulable unit is the Pod, and containers run inside Pods. Therefore, the practical way to list what’s running in a namespace is to list the Pods in that namespace. kubectl get pods shows Pods and their readiness, status, restarts, and age—giving you the canonical view of running workloads.

If you need the container-level details (images, container names), you typically use additional commands and output formatting:

kubectl describe pod <pod> to view container specs, images, states, and events

kubectl get pods -o jsonpath=... or -o wide to surface more fields

kubectl get pods -o=json to inspect .spec.containers and .status.containerStatuses

But among the provided options, kubectl get pods is the only real kubectl command that lists the running workload objects in the current namespace.

The other options are not valid kubectl subcommands: kubectl ls, kubectl ps, and kubectl show pods are not standard Kubernetes CLI operations. Kubernetes intentionally centers around the API resource model, so listing resources uses kubectl get <resource>. This also aligns with Kubernetes’ declarative nature: you observe and manage the state via API objects, not by directly enumerating OS-level processes.

So while the question says “running containers,” the Kubernetes-correct interpretation is “containers in running Pods,” and the appropriate listing command in the namespace is kubectl get pods, option A.

Question: 160

Which of the following is a good habit for cloud native cost efficiency?

- A. Follow an automated approach to cost optimization, including visibility and forecasting.
- B. Follow manual processes for cost analysis, including visibility and forecasting.
- C. Use only one cloud provider to simplify the cost analysis.
- D. Keep your legacy workloads unchanged, to avoid cloud costs.

Answer: A

Explanation:

The correct answer is A. In cloud-native environments, costs are highly dynamic: autoscaling changes compute footprint, ephemeral environments come and go, and usage-based billing applies to storage, network egress, load balancers, and observability tooling. Because of this variability, automation is the most sustainable way to achieve cost efficiency. Automated visibility (dashboards, chargeback/showback), anomaly detection, and forecasting help teams understand where spend is coming from and how it changes over time. Automated optimization actions can include right-sizing requests/limits, enforcing TTLs on preview environments, scaling down idle clusters, and cleaning unused resources.

Manual processes (B) don't scale as complexity grows. By the time someone reviews a spreadsheet or dashboard weekly, cost spikes may have already occurred. Automation enables fast feedback loops and guardrails, which is essential for preventing runaway spend caused by misconfiguration (e.g., excessive log ingestion, unbounded autoscaling, oversized node pools).

Option C is not a cost-efficiency "habit." Single-provider strategies may simplify some billing views, but they can also reduce leverage and may not be feasible for resilience/compliance; it's a business choice, not a best practice for cloud-native cost management. Option D is counterproductive: keeping legacy workloads unchanged often wastes money because cloud efficiency typically requires adapting workloads—right-sizing, adopting autoscaling, and using managed services appropriately.

In Kubernetes specifically, cost efficiency is tightly linked to resource management: accurate CPU/memory requests, limits where appropriate, cluster autoscaler tuning, and avoiding overprovisioning. Observability also matters because you can't optimize what you can't measure. Therefore, the best habit is an automated cost optimization approach with strong visibility and forecasting—A.

Question: 161

Which of these is a valid container restart policy?

- A. On login
- B. On update
- C. On start
- D. On failure

Answer: D

Explanation:

The correct answer is D: On failure. In Kubernetes, restart behavior is controlled by the Pod-level field `spec.restartPolicy`, with valid values `Always`, `OnFailure`, and `Never`. The option presented here (“On failure”) maps to Kubernetes’ `OnFailure` policy. This setting determines what the kubelet should do when containers exit:

`Always`: restart containers whenever they exit (typical for long-running services)

`OnFailure`: restart containers only if they exit with a non-zero status (common for batch workloads)

`Never`: do not restart containers (fail and leave it terminated)

So “On failure” is a valid restart policy concept and the only one in the list that matches Kubernetes semantics.

The other options are not Kubernetes restart policies. “On login,” “On update,” and “On start” are not recognized values and don’t align with how Kubernetes models container lifecycle. Kubernetes is declarative and event-driven: it reacts to container exit codes and controller intent, not user “logins.”

Operationally, choosing the right restart policy is important. For example, Jobs typically use `restartPolicy: OnFailure` or `Never` because the goal is completion, not continuous uptime. Deployments usually imply “Always” because the workload should keep serving traffic, and a crashed container should be restarted. Also note that controllers interact with restarts: a Deployment may recreate Pods if they fail readiness, while a Job counts completions and failures based on Pod termination behavior.

Therefore, among the options, the only valid (Kubernetes-aligned) restart policy is D.

Question: 162

Which of the following is a challenge derived from running cloud native applications?

- A. The operational costs of maintaining the data center of the company.

- B. Cost optimization is complex to maintain across different public cloud environments.
- C. The lack of different container images available in public image repositories.
- D. The lack of services provided by the most common public clouds.

Answer: B

Explanation:

The correct answer is B. Cloud-native applications often run across multiple environments—different cloud providers, regions, accounts/projects, and sometimes hybrid deployments. This introduces real cost-management complexity: pricing models differ (compute types, storage tiers, network egress), discount mechanisms vary (reserved capacity, savings plans), and telemetry/charge attribution can be inconsistent. When you add Kubernetes, the abstraction layer can further obscure cost drivers because costs are incurred at the infrastructure level (nodes, disks, load balancers) while consumption happens at the workload level (namespaces, Pods, services).

Option A is less relevant because cloud-native adoption often reduces dependence on maintaining a private datacenter; many organizations adopt cloud-native specifically to avoid datacenter CapEx/ops overhead. Option C is generally untrue—public registries and vendor registries contain vast numbers of images; the challenge is more about provenance, security, and supply chain than “lack of images.” Option D is incorrect because major clouds offer abundant services; the difficulty is choosing among them and controlling cost/complexity, not a lack of services.

Cost optimization being complex is a recognized challenge because cloud-native architectures include microservices sprawl, autoscaling, ephemeral environments, and pay-per-use dependencies (managed databases, message queues, observability). Small misconfigurations can cause big bills: noisy logs, over-requested resources, unbounded HPA scaling, and egress-heavy architectures. That’s why practices like FinOps, tagging/labeling for allocation, and automated guardrails are emphasized.

So the best answer describing a real, common cloud-native challenge is B.

Question: 163

What is the correct hierarchy of Kubernetes components?

- A. Containers → Pods → Cluster → Nodes
- B. Nodes → Cluster → Containers → Pods

C. Cluster → Nodes → Pods → Containers

D. Pods → Cluster → Containers → Nodes

Answer: C

Explanation:

The correct answer is C: Cluster → Nodes → Pods → Containers. This expresses the fundamental structural relationship in Kubernetes. A cluster is the overall system (control plane + nodes) that runs your workloads. Inside the cluster, you have nodes (worker machines—VMs or bare metal) that provide CPU, memory, storage, and networking. The scheduler assigns workloads to nodes.

Workloads are executed as Pods, which are the smallest deployable units Kubernetes schedules. Pods represent one or more containers that share networking (one Pod IP and port space) and can share storage volumes. Within each Pod are containers, which are the actual application processes packaged with their filesystem and runtime dependencies.

The other options are incorrect because they break these containment relationships. Containers do not contain Pods; Pods contain containers. Nodes do not exist "inside" Pods; Pods run on nodes. And the cluster is the top-level boundary that contains nodes and orchestrates Pods.

This hierarchy matters for troubleshooting and design. If you're thinking about capacity, you reason at the node and cluster level (node pools, autoscaling, quotas). If you're thinking about application scaling, you reason at the Pod level (replicas, HPA, readiness probes). If you're thinking about process-level concerns, you reason at the container level (images, security context, runtime user, resources). Kubernetes intentionally uses this layered model so that scheduling and orchestration operate on Pods, while the container runtime handles container execution details.

So the accurate hierarchy from largest to smallest unit is: Cluster → Nodes → Pods → Containers, which corresponds to C.

Question: 164

Which statement about Secrets is correct?

A. A Secret is part of a Pod specification.

B. Secret data is encrypted with the cluster private key by default.

C. Secret data is base64 encoded and stored unencrypted by default.

D. A Secret can only be used for confidential data.

Answer: C

Explanation:

The correct answer is C. By default, Kubernetes Secrets store their data as base64-encoded values in the API (backed by etcd). Base64 is an encoding mechanism, not encryption, so this does not provide confidentiality. Unless you explicitly configure encryption at rest for etcd (via the API server encryption provider configuration) and secure access controls, Secret contents should be treated as potentially readable by anyone with sufficient API access or access to etcd backups.

Option A is misleading: a Secret is its own Kubernetes resource (kind: Secret). While Pods can reference Secrets (as environment variables or mounted volumes), the Secret itself is not "part of the Pod spec" as an embedded object.

Option B is incorrect because Kubernetes does not automatically encrypt Secret data with a cluster private key by default; encryption at rest is optional and must be enabled. Option D is incorrect because Secrets can store a range of sensitive or semi-sensitive data (tokens, certs, passwords), but Kubernetes does not enforce "only confidential data" semantics; it's a storage mechanism with size and format constraints.

Operationally, best practices include: enabling encryption at rest, limiting access via RBAC, avoiding broad "list/get secrets" permissions, using dedicated service accounts, auditing access, and considering external secrets managers (Vault, cloud KMS-backed solutions) for higher assurance. Also, don't confuse "Secret" with "secure by default." The default protection is mainly about avoiding accidental plaintext exposure in manifests, not about cryptographic security.

So the only correct statement in the options is C.

Question: 165

Which mechanism allows extending the Kubernetes API?

- A. ConfigMap
- B. CustomResourceDefinition
- C. MutatingAdmissionWebhook mechanism
- D. Kustomize

Answer: B

Explanation:

The correct answer is B: CustomResourceDefinition (CRD). Kubernetes is designed to be extensible. A CRD lets you define your own resource types (custom API objects) that behave like native Kubernetes resources: they can be created with YAML, stored in etcd, retrieved via the API server, and managed using kubectl. For example, operators commonly define CRDs such as Databases, RedisClusters, or Certificates to model higher-level application concepts.

A CRD extends the API by adding a new kind under a group/version (e.g., example.com/v1). You typically pair CRDs with a controller (often called an operator) that watches these custom objects and reconciles real-world resources (Deployments, StatefulSets, cloud resources) to match the desired state specified in the CRD instances. This is the same control-loop pattern used for built-in controllers—just applied to your custom domain.

Why the other options aren't correct: ConfigMaps store configuration data but do not add new API types. A MutatingAdmissionWebhook can modify or validate requests for existing resources, but it doesn't define new API kinds; it enforces policy or injects defaults. Kustomize is a manifest customization tool (patch/overlay) and doesn't extend the Kubernetes API surface.

CRDs are foundational to much of the Kubernetes ecosystem: cert-manager, Argo, Istio, and many operators rely heavily on CRDs. They also support schema validation via OpenAPI v3 schemas, which improves safety and tooling (better error messages, IDE hints). Therefore, the mechanism for

extending the Kubernetes API is CustomResourceDefinition, option B.

Question: 166

Which of the following observability data streams would be most useful when desiring to plot resource consumption and predicted future resource exhaustion?

- A. stdout
- B. Traces
- C. Logs
- D. Metrics

Answer: D

Explanation:

The correct answer is D: Metrics. Metrics are numeric time-series measurements collected at regular intervals, making them ideal for plotting resource consumption over time and forecasting future exhaustion. In Kubernetes, this includes CPU usage, memory usage, disk I/O, network throughput, filesystem usage, Pod restarts, and node allocatable vs requested resources. Because metrics are structured and queryable (often with Prometheus), you can compute rates, aggregates, percentiles, and trends, and then apply forecasting methods to predict when a resource will run out.

Logs and traces have different purposes. Logs are event records (strings) that are great for debugging and auditing, but they are not naturally suited to continuous quantitative plotting unless you transform them into metrics (log-based metrics). Traces capture end-to-end request paths and latency breakdowns; they help you find slow spans and dependency bottlenecks, not forecast CPU/memory exhaustion. stdout is just a stream where logs might be written; by itself it's not an observability data type used for capacity trending.

In Kubernetes observability stacks, metrics are typically scraped from components and workloads: kubelet/cAdvisor exports container metrics, node exporters expose host metrics, and applications expose business/system metrics. The metrics pipeline (Prometheus, OpenTelemetry metrics, managed monitoring) enables dashboards and alerting. For resource exhaustion, you often alert on "time to fill" (e.g., predicted disk fill in < N hours), high sustained utilization, or rapidly increasing error rates due to throttling.

Therefore, the most appropriate data stream for plotting consumption and predicting exhaustion is Metrics, option D.

Question: 167

What can be used to create a job that will run at specified times/dates or on a repeating schedule?

- A. Job
- B. CalendarJob
- C. BatchJob
- D. CronJob

Answer: D

Explanation:

The correct answer is D: CronJob. A Kubernetes CronJob is specifically designed for creating Jobs on a schedule—either at specified times/dates (expressed via cron syntax) or on a repeating interval (hourly, daily, weekly). When the schedule triggers, the CronJob controller creates a Job, and the Job controller creates the Pods that execute the workload to completion.

Option A (Job) is not inherently scheduled. A Job runs when you create it, and it continues until it completes successfully or fails according to its retry/backoff behavior. If you want it to run periodically, you need something else to create the Job each time. CronJob is the built-in mechanism for that scheduling.

Options B and C are not standard Kubernetes workload objects. Kubernetes does not include “CalendarJob” or “BatchJob” as official API kinds. The scheduling primitive is CronJob.

CronJobs also include important operational controls: concurrency policies prevent overlapping runs, deadlines control missed schedules, and history limits manage old Job retention. This makes CronJobs more robust than ad-hoc scheduling approaches and keeps the workload lifecycle visible in the Kubernetes API (status/events/logs). It also means you can apply standard Kubernetes patterns: use a service account with least privilege, mount Secrets/ConfigMaps, run in specific namespaces, and manage resource requests/limits so that scheduled workloads don't destabilize the cluster.

So the correct Kubernetes resource for scheduled and repeating job execution is CronJob (D).

Question: 168

If kubectl is failing to retrieve information from the cluster, where can you find Pod logs to troubleshoot?

- A. /var/log/pods/
- B. ~/.kube/config
- C. /var/log/k8s/
- D. /etc/kubernetes/

Answer: A

Explanation:

The correct answer is A: `/var/log/pods/`. When `kubectl` logs can't retrieve logs (for example, API connectivity issues, auth problems, or kubelet/API proxy issues), you can often troubleshoot directly on the node where the Pod ran. Kubernetes nodes typically store container logs on disk, and a common location is under `/var/log/pods/`, organized by namespace, Pod name/UID, and container. This directory contains symlinks or files that map to the underlying container runtime log location (often under `/var/log/containers/` as well, depending on distro/runtime setup).

Option B (`~/.kube/config`) is your local kubeconfig file; it contains cluster endpoints and credentials, not Pod logs. Option D (`/etc/kubernetes/`) contains Kubernetes component configuration/manifests on some installations (especially control plane), not application logs. Option C (`/var/log/k8s/`) is not a standard Kubernetes log path.

Operationally, the node-level log locations depend on the container runtime and logging configuration, but the Kubernetes convention is that kubelet writes container logs to a known location and exposes them through the API so `kubectl` logs works. If the API path is broken, node access becomes your fallback. This is also why secure node access is sensitive: anyone with node root access can potentially read logs (and other data), which is part of the threat model.

So, the best answer for where to look on the node for Pod logs when `kubectl` can't retrieve them is `/var/log/pods/`, option A.

Question: 169

Which component in Kubernetes is responsible to watch newly created Pods with no assigned node, and selects a node for them to run on?

- A. etcd
- B. kube-controller-manager
- C. kube-proxy
- D. kube-scheduler

Answer: D

Explanation:

The correct answer is D: kube-scheduler. The kube-scheduler is the control plane component responsible for assigning Pods to nodes. It watches for newly created Pods that do not have a `spec.nodeName` set (i.e., unscheduled Pods). For each such Pod, it evaluates the available nodes against scheduling constraints and chooses the best node, then performs a "bind" operation by setting the Pod's `spec.nodeName`.

Scheduling decisions consider many factors: resource requests vs node allocatable capacity, taints/tolerations, node selectors and affinity/anti-affinity, topology spread constraints, and other

policy inputs. The scheduler typically runs a two-phase process: filtering (find feasible nodes) and scoring (rank feasible nodes) before selecting one.

Option A (etcd) is the datastore that persists cluster state; it does not make scheduling decisions. Option B (kube-controller-manager) runs controllers (Deployment, Node, Job controllers, etc.) but not scheduling. Option C (kube-proxy) is a node component for Service networking; it doesn't place Pods.

Understanding this separation is key for troubleshooting. If Pods are stuck Pending with "no nodes available," the scheduler's feasibility checks are failing (insufficient CPU/memory, taints not tolerated, affinity mismatch). If Pods schedule but land unexpectedly, it's often due to scoring preferences or missing constraints. In all cases, the component that performs the node selection is the kube-scheduler.

Therefore, the verified correct answer is D.

Question: 170

Which control plane component is responsible for updating the node Ready condition if a node becomes unreachable?

- A. The kube-proxy
- B. The node controller
- C. The kubect
- D. The kube-apiserver

Answer: B

Explanation:

The correct answer is B: the node controller. In Kubernetes, node health is monitored and reflected through Node conditions such as Ready. The Node Controller (a controller that runs as part of the control plane, within the controller-manager) is responsible for monitoring node heartbeats and

updating node status when a node becomes unreachable or unhealthy.

Nodes periodically report status (including kubelet heartbeats) to the API server. The Node Controller watches these updates. If it detects that a node has stopped reporting within expected time windows, it marks the node condition Ready as Unknown (or otherwise updates conditions) to indicate the control plane can't confirm node health. This status change then influences higher-level behaviors such as Pod eviction and rescheduling: after grace periods and eviction timeouts, Pods on an unhealthy node may be evicted so the workload can be recreated on healthy nodes (assuming a controller manages replicas).

Option A (kube-proxy) is a node component for Service traffic routing and does not manage node health conditions. Option C (kubectl) is a CLI client; it does not participate in control plane health monitoring. Option D (kube-apiserver) stores and serves Node status, but it doesn't decide when a node is unreachable; it persists what controllers and kubelets report. The "decision logic" for updating the Ready condition in response to missing heartbeats is the Node Controller's job.

So, the component that updates the Node Ready condition when a node becomes unreachable is the **node controller**, which is option B.

Question: 171

Which GitOps engine can be used to orchestrate parallel jobs on Kubernetes?

- A. Jenkins X
- B. Flagger
- C. Flux
- D. Argo Workflows

Answer: D

Explanation:

Argo Workflows (D) is the correct answer because it is a Kubernetes-native workflow engine

designed to define and run multi-step workflows—often with parallelization—directly on Kubernetes. Argo Workflows models workflows as DAGs (directed acyclic graphs) or step-based sequences, where each step is typically a Pod. Because each step is expressed as Kubernetes resources (custom resources), Argo can schedule many tasks concurrently, control fan-out/fan-in patterns, and manage dependencies between steps (e.g., “run these 10 jobs in parallel, then aggregate results”).

The question calls it a “GitOps engine,” but the capability being tested is “orchestrate parallel jobs.” Argo Workflows fits because it is purpose-built for running complex job orchestration, including parallel tasks, retries, timeouts, artifacts passing, and conditional execution. In practice, many teams store workflow manifests in Git and apply GitOps practices around them, but the distinguishing feature here is the workflow orchestration engine itself.

Why the other options are not best:

Flux (C) is a GitOps controller that reconciles cluster state from Git; it doesn’t orchestrate parallel job graphs as its core function.

Flagger (B) is a progressive delivery operator (canary/blue-green) often paired with GitOps and service meshes/Ingress; it’s not a general workflow orchestrator for parallel batch jobs.

Jenkins X (A) is CI/CD-focused (pipelines), not primarily a Kubernetes-native workflow engine for parallel job DAGs in the way Argo Workflows is.

So, the Kubernetes-native tool specifically used to orchestrate parallel jobs and workflows is Argo Workflows (D).

Question: 172

What is the main purpose of the Open Container Initiative (OCI)?

A. Accelerating the adoption of containers and Kubernetes in the industry.

B. Creating open industry standards around container formats and runtimes.

C. Creating industry standards around container formats and runtimes for private purposes.

D. Improving the security of standards around container formats and runtimes.

Answer: B

Explanation:

B is correct: the OCI's main purpose is to create open, vendor-neutral industry standards for container image formats and container runtimes. Standardization is critical in container orchestration because portability is a core promise: you should be able to build an image once and run it across different environments and runtimes without rewriting packaging or execution logic.

OCI defines (at a high level) two foundational specs:

Image specification: how container images are packaged (layers, metadata, manifests).

Runtime specification: how to run a container (filesystem setup, namespaces/cgroups behavior, lifecycle).

These standards enable interoperability across tooling. For example, higher-level runtimes (like containerd or CRI-O) rely on OCI-compliant components (often runc or equivalents) to execute containers consistently.

Why the other options are not the best answer:

A (accelerating adoption) might be an indirect outcome, but it's not the OCI's core charter.

C is contradictory ("industry standards" but "for private purposes")—OCI is explicitly about open standards.

D (improving security) can be helped by standardization and best practices, but OCI is not primarily a security standards body; its central function is format and runtime interoperability.

In Kubernetes specifically, OCI is part of the "plumbing" that makes runtimes replaceable.

Kubernetes talks to runtimes via CRI; runtimes execute containers via OCI. This layering helps Kubernetes remain runtime-agnostic while still benefiting from consistent container behavior everywhere.

Therefore, the correct choice is B: OCI creates open standards around container formats and runtimes.

Question: 173

Which are the core features provided by a service mesh?

A. Authentication and authorization

B. Distributing and replicating data

- C. Security vulnerability scanning
- D. Configuration management

Answer: A

Explanation:

A is the correct answer because a service mesh primarily focuses on securing and managing service- to-service communication, and a core part of that is authentication and authorization. In microservices architectures, internal (“east-west”) traffic can become a complex web of calls. A service mesh introduces a dedicated communication layer—commonly implemented with sidecar proxies or node proxies plus a control plane—to apply consistent security and traffic policies across Services.

Authentication in a mesh typically means service identity: each workload gets an identity (often via certificates), enabling mutual TLS (mTLS) so services can verify each other and encrypt traffic in transit. Authorization then builds on identity to enforce “who can talk to whom” via policies (for example: service A can call service B only on certain paths or methods). These capabilities are central because they reduce the need for every development team to implement and maintain custom security libraries correctly.

Why the other answers are incorrect:

B (data distribution/replication) is a storage/database concern, not a mesh function.

C (vulnerability scanning) is typically part of CI/CD and supply-chain security tooling, not service-to- service runtime traffic management.

D (configuration management) is broader (GitOps, IaC, Helm/Kustomize); a mesh does have configuration, but “configuration management” is not the defining core feature tested here.

Service meshes also commonly provide traffic management (timeouts, retries, circuit breaking, canary routing) and telemetry (metrics/traces), but among the listed options, authentication and authorization best matches “core features.” It captures the mesh’s role in standardizing secure communications in a distributed system.

So, the verified correct answer is A.

Question: 174

Which of the following options include only mandatory fields to create a Kubernetes object using a YAML file?

- A. apiVersion, template, kind, status
- B. apiVersion, metadata, status, spec
- C. apiVersion, template, kind, spec
- D. apiVersion, metadata, kind, spec

Answer: D

Explanation:

D is correct: the mandatory top-level fields for creating a Kubernetes object manifest are apiVersion, kind, metadata, and (for most objects you create) spec. These fields establish what the object is and what you want Kubernetes to do with it.

apiVersion tells Kubernetes which API group/version schema to use (e.g., apps/v1, v1). This determines valid fields and behavior.

kind identifies the resource type (e.g., Pod, Deployment, Service).

metadata contains identifying information like name, namespace, and labels/annotations used for organization, selection, and automation.

spec describes the desired state. Controllers and the kubelet reconcile actual state to match spec.

Why other choices are wrong:

status is not a mandatory input field. It's generally written by Kubernetes controllers and reflects observed state (conditions, readiness, assigned node, etc.). Users typically do not set status when creating objects.

template is not a universal top-level field. It exists inside some resources (notably Deployment.spec.template), but it's not a required top-level field across Kubernetes objects.

It's true that some resources can be created without a spec (or with minimal fields), but in the examstyle framing—"mandatory fields... using a YAML file"—the canonical expected set is exactly the four in D. This aligns with how Kubernetes documentation and examples present manifests: identify the API schema and kind, give object

metadata, and declare desired state.

Therefore, apiVersion + metadata + kind + spec is the only option that includes only the mandatory fields, making D the verified correct answer.

Question: 175

Which of the following is the name of a container orchestration software?

- A. OpenStack
- B. Docker
- C. Apache Mesos
- D. CRI-O

Answer: C

Explanation:

C (Apache Mesos) is correct because Mesos is a cluster manager/orchestrator that can schedule and manage workloads (including containerized workloads) across a pool of machines. Historically, Mesos (often paired with frameworks like Marathon) was used to orchestrate services and batch jobs at scale, similar in spirit to Kubernetes' scheduling and cluster management role.

Why the other answers are not correct as "container orchestration software" in this context:

OpenStack (A) is primarily an IaaS cloud platform for provisioning compute, networking, and storage (VM-focused). It's not a container orchestrator, though it can host Kubernetes or containers.

Docker (B) is a container platform/tooling ecosystem (image build, runtime, local orchestration via Docker Compose/Swarm historically), but "Docker" itself is not the best match for "container orchestration software" in the multi-node cluster orchestration sense that the question implies.

CRI-O (D) is a container runtime implementing Kubernetes' CRI; it runs containers on a node but does not orchestrate placement, scaling, or service lifecycle across a cluster.

Container orchestration typically means capabilities like scheduling, scaling, service discovery integration, health management, and rolling updates across multiple hosts. Mesos fits that definition: it provides resource management and scheduling over a cluster and can run container workloads via supported containerizers. Kubernetes ultimately became the dominant orchestrator for many use cases, but Mesos is clearly recognized as orchestration software in this category. So, among these choices, the verified orchestration platform is Apache Mesos (C).

Question: 176

What happens with a regular Pod running in Kubernetes when a node fails?

- A. A new Pod with the same UID is scheduled to another node after a while.
- B. A new, near-identical Pod but with different UID is scheduled to another node.
- C. By default, a Pod can only be scheduled to the same node when the node fails.
- D. A new Pod is scheduled on a different node only if it is configured explicitly.

Answer: B

Explanation:

B is correct: when a node fails, Kubernetes does not “move” the same Pod instance; instead, a new

Pod object (new UID) is created to replace it—assuming the Pod is managed by a controller (Deployment/ReplicaSet, StatefulSet, etc.). A Pod is an API object with a unique identifier (UID) and is tightly associated with the node it's scheduled to via `spec.nodeName`. If the node becomes unreachable, that original Pod cannot be restarted elsewhere because it was bound to that node.

Kubernetes' high availability comes from controllers maintaining desired state. For example, a Deployment desires `N` replicas. If a node fails and the replicas on that node are lost, the controller will create replacement Pods, and the scheduler will place them onto healthy nodes. These replacement Pods will be “near-identical” in spec (same template), but they are still new instances with new UIDs and typically new IPs.

Why the other options are wrong:

A is incorrect because the UID does not remain the same—Kubernetes creates a new Pod object rather than reusing

the old identity.

C is incorrect; pods are not restricted to the same node after failure. The whole point of orchestration is to reschedule elsewhere.

D is incorrect; rescheduling does not require special explicit configuration for typical controller-managed workloads. The controller behavior is standard. (If it's a bare Pod without a controller, it will not be recreated automatically.)

This also ties to the difference between "regular Pod" vs controller-managed workloads: a standalone Pod is not self-healing by itself, while a Deployment/ReplicaSet provides that resilience. In typical production design, you run workloads under controllers specifically so node failure triggers replacement and restores replica count.

Therefore, the correct outcome is B.

Question: 177

What is the minimum number of etcd members that are required for a highly available Kubernetes cluster?

- A. Two etcd members.
- B. Five etcd members.
- C. Six etcd members.
- D. Three etcd members.

Answer: D

Explanation:

D (three etcd members) is correct. etcd is a distributed key-value store that uses the Raft consensus algorithm. High availability in consensus systems depends on maintaining a quorum (majority) of members to continue serving writes reliably. With 3 members, the cluster can tolerate 1 failure and still have 2/3 available—enough for quorum.

Two members is a common trap: with 2, a single failure leaves 1/2, which is not a majority, so the cluster cannot safely make progress. That means 2-member etcd is not HA; it is fragile and can be taken down by one node loss, network partition, or maintenance event. Five members can tolerate 2 failures and is a valid HA configuration, but it is not the minimum. Six is even-sized and generally discouraged for consensus because it doesn't improve failure

tolerance compared to five (quorum still requires 4), while increasing coordination overhead.

In Kubernetes, etcd reliability directly affects the API server and the entire control plane because etcd stores cluster state: object specs, status, controller state, and more. If etcd loses quorum, the API server will be unable to persist or reliably read/write state, leading to cluster management outages. That's why the minimum HA baseline is three etcd members, often across distinct failure domains (nodes/AZs), with strong disk performance and consistent low-latency networking.

So, the smallest etcd topology that provides true fault tolerance is 3 members, which corresponds to option D.

Question: 178

What is the main purpose of the Ingress in Kubernetes?

- A. Access HTTP and HTTPS services running in the cluster based on their IP address.
- B. Access services different from HTTP or HTTPS running in the cluster based on their IP address.
- C. Access services different from HTTP or HTTPS running in the cluster based on their path.
- D. Access HTTP and HTTPS services running in the cluster based on their path.

Answer: D

Explanation:

D is correct. Ingress is a Kubernetes API object that defines rules for external access to HTTP/HTTPS services in a cluster. The defining capability is Layer 7 routing—commonly host-based and path-based routing—so you can route requests like `example.com/app1` to one Service and `example.com/app2` to another. While the question mentions “based on their path,” that’s a classic and correct Ingress use case (and host routing is also common).

Ingress itself is only the specification of routing rules. An Ingress controller (e.g., NGINX Ingress Controller, HAProxy, Traefik, cloud-provider controllers) is what actually implements those rules by configuring a reverse proxy/load balancer. Ingress typically terminates TLS (HTTPS) and forwards traffic to internal Services, giving a more expressive alternative to exposing every service via NodePort/LoadBalancer.

Why the other options are wrong:

A suggests routing by IP address; Ingress is fundamentally about HTTP(S) routing rules (host/path), not direct Service IP access.

B and C describe non-HTTP protocols; Ingress is specifically for HTTP/HTTPS. For TCP/UDP or other protocols, you generally use Services of type LoadBalancer/NodePort, Gateway API implementations, or controller-specific TCP/UDP configuration.

Ingress is a foundational building block for cloud-native application delivery because it centralizes edge routing, enables TLS management, and supports gradual adoption patterns (multiple services under one domain). Therefore, the main purpose described here matches D.

Question: 179

How can you extend the Kubernetes API?

- A. Adding a CustomResourceDefinition or implementing an aggregation layer.
- B. Adding a new version of a resource, for instance v4beta3.
- C. With the command `kubectl extend api`, logged in as an administrator.
- D. Adding the desired API object as a kubelet parameter.

Answer: A

Explanation:

A is correct: Kubernetes' API can be extended by adding CustomResourceDefinitions (CRDs) and/or by implementing the API Aggregation Layer. These are the two canonical extension mechanisms.

CRDs let you define new resource types (new kinds) that the Kubernetes API server stores in etcd and serves like native objects. You typically pair a CRD with a controller/operator that watches those custom objects and reconciles real resources accordingly. This pattern is foundational to the Kubernetes ecosystem (many popular add-ons install CRDs).

The aggregation layer allows you to add entire API services (aggregated API servers) that serve additional endpoints under the Kubernetes API. This is used when you want custom API behavior, custom storage, or specialized semantics beyond what CRDs provide (or when implementing APIs like metrics APIs historically).

Why the other answers are wrong:

B is not how API extension works. You don't "extend the API" by inventing new versions like v4beta3; versions are defined and implemented by API servers/controllers, not by users arbitrarily.

C is fictional; there is no standard kubectl extend api command.

D is also incorrect; kubelet parameters configure node agent behavior, not API server types and discovery.

So, the verified ways to extend Kubernetes' API surface are CRDs and API aggregation, which is option A.

Question: 180

What is an ephemeral container?

- A. A specialized container that runs as root for infosec applications.
- B. A specialized container that runs temporarily in an existing Pod.
- C. A specialized container that extends and enhances the main container in a Pod.
- D. A specialized container that runs before the app container in a Pod.

Answer: B

Explanation:

B is correct: an ephemeral container is a temporary container you can add to an existing Pod for troubleshooting and debugging without restarting the Pod. This capability is especially useful when a running container image is minimal (distroless) and lacks debugging tools like sh, curl, or ps. Instead of rebuilding the workload image or disrupting the Pod, you attach an ephemeral container that includes the tools you need, then inspect processes, networking, filesystem mounts, and runtime behavior.

Ephemeral containers are not part of the original Pod spec the same way normal containers are. They are added via a dedicated subresource and are generally not restarted automatically like regular containers. They are meant for interactive investigation, not for ongoing workload functionality.

Why the other options are incorrect:

D describes init containers, which run before app containers start and are used for setup tasks.

C resembles the “sidecar” concept (a supporting container that runs alongside the main container), but sidecars are normal containers defined in the Pod spec, not ephemeral containers.

A is not a definition; ephemeral containers are not “root by design” (they can run with various security contexts depending on policy), and they aren’t limited to infosec use cases.

In Kubernetes operations, ephemeral containers complement kubectl exec and logs. If the target container is crash-looping or lacks a shell, exec may not help; adding an ephemeral container provides a safe and Kubernetes-native debugging path. So, the accurate definition is B.

Question: 181

In a cloud native environment, who is usually responsible for maintaining the workloads running across the different platforms?

- A. The cloud provider.
- B. The Site Reliability Engineering (SRE) team.
- C. The team of developers.
- D. The Support Engineering team (SE).

Answer: B

Explanation:

B (the Site Reliability Engineering team) is correct. In cloud-native organizations, SREs are commonly responsible for the reliability, availability, and operational health of workloads across platforms (multiple clusters, regions, clouds, and supporting services). While responsibilities vary by company, the classic SRE charter is to apply software engineering to operations: build automation, standardize runbooks, manage incident response, define SLOs/SLIs, and continuously improve system reliability.

Maintaining workloads “across different platforms” implies cross-cutting operational ownership: deployments need to behave consistently, rollouts must be safe, monitoring and alerting must be uniform, and incident practices must work across environments. SRE teams typically own or heavily influence the observability stack (metrics/logs/traces), operational readiness, capacity planning, and reliability guardrails (error budgets, progressive delivery, automated rollback triggers). They also collaborate closely with platform engineering and application teams, but SRE is often the group that ensures production workloads meet reliability targets.

Why other options are less correct:

The cloud provider (A) maintains the underlying cloud services, but not your application workloads' correctness, SLOs, or operational processes.

Developers (C) do maintain application code and may own on-call in some models, but the question asks "usually" in cloud-native environments; SRE is the widely recognized function for workload reliability across platforms.

Support Engineering (D) typically focuses on customer support and troubleshooting from a user perspective, not maintaining platform workload reliability at scale.

So, the best and verified answer is B: SRE teams commonly maintain and ensure reliability of workloads across cloud-native platforms.

Question: 182

Which Kubernetes-native deployment strategy supports zero-downtime updates of a workload?

- A. Canary
- B. Recreate
- C. BlueGreen
- D. RollingUpdate

Answer: D

Explanation:

D (RollingUpdate) is correct. In Kubernetes, the Deployment resource's default update strategy is RollingUpdate, which replaces Pods gradually rather than all at once. This supports zero-downtime updates when the workload is properly configured (sufficient replicas, correct readiness probes, and appropriate maxUnavailable / maxSurge settings). As new Pods come up and become Ready, old Pods are terminated in a controlled way, keeping the service available throughout the rollout.

RollingUpdate's "zero downtime" is achieved by maintaining capacity while transitioning between versions. For example, with multiple replicas, Kubernetes can create new Pods, wait for readiness, then scale down old Pods, ensuring traffic continues to flow to healthy instances. Readiness probes are critical: they prevent traffic from being routed to a Pod until

it's actually ready to serve.

Why other options are not the Kubernetes-native "strategy" answer here:

Recreate (B) explicitly stops old Pods before starting new ones, causing downtime for most services.

Canary (A) and BlueGreen (C) are real deployment patterns, but in "Kubernetes-native deployment strategy" terms, the built-in Deployment strategies are RollingUpdate and Recreate.

Canary/BlueGreen typically require additional tooling/controllers (service mesh, ingress controller features, or progressive delivery operators) to manage traffic shifting between versions.

So, for a Kubernetes-native strategy that supports zero-downtime updates, the correct and verified choice is RollingUpdate (D).

Question: 183

What service account does a Pod use in a given namespace when the service account is not specified?

- A. admin
- B. sysadmin
- C. root
- D. default

Answer: D

Explanation:

D (default) is correct. In Kubernetes, if you create a Pod (or a controller creates Pods) without specifying `spec.serviceAccountName`, Kubernetes assigns the Pod the default ServiceAccount in that namespace. The ServiceAccount determines what identity the Pod uses when accessing the Kubernetes API (for example, via the in-cluster token mounted into the Pod, when token automounting is enabled).

Every namespace typically has a default ServiceAccount created automatically. The permissions associated with that

ServiceAccount are determined by RBAC bindings. In many clusters, the default ServiceAccount has minimal permissions (or none) as a security best practice, because leaving it overly privileged would allow any Pod to access sensitive cluster APIs.

Why the other options are wrong: Kubernetes does not automatically choose "admin," "sysadmin," or "root" service accounts. Those are not standard implicit identities, and automatically granting admin privileges would be insecure. Instead, Kubernetes follows a predictable, least-privilege- friendly default: use the namespace's default ServiceAccount unless you explicitly request a different one.

Operationally, this matters for security and troubleshooting. If an application in a Pod is failing with "forbidden" errors when calling the API, it often means it's using the default ServiceAccount without the necessary RBAC permissions. The correct fix is usually to create a dedicated ServiceAccount and bind only the required roles, then set serviceAccountName in the Pod template. Conversely, if you're hardening a cluster, you often disable automounting of service account tokens for Pods that don't need API access.

Therefore, the verified correct answer is D: default.

Question: 184

What is a cloud native application?

- A. It is a monolithic application that has been containerized and is running now on the cloud.
- B. It is an application designed to be scalable and take advantage of services running on the cloud.
- C. It is an application designed to run all its functions in separate containers.
- D. It is any application that runs in a cloud provider and uses its services.

Answer: B

Explanation:

B is correct. A cloud native application is designed to be scalable, resilient, and adaptable, and to leverage cloud/platform capabilities rather than merely being "hosted" on a cloud VM. Cloud-native design emphasizes principles like elasticity (scale up/down), automation, fault tolerance, and rapid, reliable delivery. While containers and Kubernetes are common enablers, the key is the architectural intent: build applications that embrace distributed systems patterns and cloud-managed primitives.

Option A is not enough. Simply containerizing a monolith and running it in the cloud does not automatically make it cloud native; that may be “lift-and-shift” packaging. The application might still be tightly coupled, hard to scale, and operationally fragile. Option C is too narrow and prescriptive; cloud native does not require “all functions in separate containers” (microservices are common but not mandatory). Many cloud-native apps use a mix of services, and even monoliths can be made more cloud native by adopting statelessness, externalized state, and automated delivery. Option D is too broad; “any app running in a cloud provider” includes legacy apps that don’t benefit from elasticity or cloud-native operational models.

Cloud-native applications typically align with patterns: stateless service tiers, declarative configuration, health endpoints, horizontal scaling, graceful shutdown, and reliance on managed backing services (databases, queues, identity, observability). They are built to run reliably in dynamic environments where instances are replaced routinely—an assumption that matches Kubernetes’ reconciliation and self-healing model.

So, the best verified definition among these options is B.

Question: 185

What's the most adopted way of conflict resolution and decision-making for the open-source projects under the CNCF umbrella?

- A. Financial Analysis
- B. Discussion and Voting
- C. Flipism Technique
- D. Project Founder Say

Answer: B

Explanation:

B (Discussion and Voting) is correct. CNCF-hosted open-source projects generally operate with open governance practices that emphasize transparency, community participation, and documented decision-making. While each project can have its own governance model (maintainers, technical steering committees, SIGs, TOC interactions, etc.), a very common and widely adopted approach to resolving disagreements and making

decisions is to first pursue discussion (often on GitHub issues/PRs, mailing lists, or community meetings) and then use voting/consensus mechanisms when needed.

This approach is important because open-source communities are made up of diverse contributors across companies and geographies. “Project Founder Say” (D) is not a sustainable or typical CNCF governance norm for mature projects; CNCF explicitly encourages neutral, community-led governance rather than single-person control. “Financial Analysis” (A) is not a conflict resolution mechanism for technical decisions, and “Flipism Technique” (C) is not a real governance practice.

In Kubernetes specifically, community decisions are often made within structured groups (e.g., SIGs) using discussion and consensus-building, sometimes followed by formal votes where governance requires it. The goal is to ensure decisions are fair, recorded, and aligned with the project’s mission and contributor expectations. This also reduces risk of vendor capture and builds trust: anyone can review the rationale in meeting notes, issues, or PR threads, and decisions can be revisited with new evidence.

Therefore, the most adopted conflict resolution and decision-making method across CNCF opensource projects is discussion and voting, making B the verified correct answer.

Question: 186

Which of the following options include resources cleaned by the Kubernetes garbage collection mechanism?

- A. Stale or expired CertificateSigningRequests (CSRs) and old deployments.
- B. Nodes deleted by a cloud controller manager and obsolete logs from the kubelet.
- C. Unused container and container images, and obsolete logs from the kubelet.
- D. Terminated pods, completed jobs, and objects without owner references.

Answer: D

Explanation:

Kubernetes garbage collection (GC) is about cleaning up API objects and related resources that are no longer needed, so the correct answer is D. Two big categories it targets are (1) objects that have finished their lifecycle (like terminated Pods and completed Jobs, depending on controllers and TTL policies), and (2) “dangling” objects that are no longer referenced properly—often described as objects without owner references (or where owners are gone), which can

happen when a higher-level controller is deleted or when dependent resources are left behind.

A key Kubernetes concept here is OwnerReferences: many resources are created “owned” by a controller (e.g., a ReplicaSet owned by a Deployment, Pods owned by a ReplicaSet). When an owning object is deleted, Kubernetes’ garbage collector can remove dependent objects based on deletion propagation policies (foreground/background/orphan). This prevents resource leaks and keeps the cluster tidy and performant.

The other options are incorrect because they refer to cleanup tasks outside Kubernetes GC’s scope. Kubelet logs (B/C) are node-level files and log rotation is handled by node/runtime configuration, not the Kubernetes garbage collector. Unused container images (C) are managed by the container runtime’s image GC and kubelet disk pressure management, not the Kubernetes API GC. Nodes deleted by a cloud controller (B) aren’t “garbage collected” in the same sense; node lifecycle is handled by controllers and cloud integrations, but not as a generic GC cleanup category like ownerRef-based object deletion.

So, when the question asks specifically about “resources cleaned by Kubernetes garbage collection,” it’s pointing to Kubernetes object lifecycle cleanup: terminated Pods, completed Jobs, and orphaned objects—exactly what option D states.

Question: 187

What is the default eviction timeout when the Ready condition of a node is Unknown or False?

- A. Thirty seconds.
- B. Thirty minutes.
- C. One minute.
- D. Five minutes.

Answer: D

Explanation:

The verified correct answer is D (Five minutes). In Kubernetes, node health is continuously monitored. When a node stops reporting status (heartbeats from the kubelet) or is otherwise considered unreachable, the Node controller updates the Node’s Ready condition to Unknown (or it can become False). From that point, Kubernetes has to balance two risks: acting too quickly might cause unnecessary disruption (e.g., transient network hiccups), but acting too slowly prolongs

outage for workloads that were running on the failed node.

The “default eviction timeout” refers to the control plane behavior that determines how long Kubernetes waits before evicting Pods from a node that appears unhealthy/unreachable. After this timeout elapses, Kubernetes begins eviction of Pods so controllers (like Deployments) can recreate them on healthy nodes, restoring the desired replica count and availability.

This is tightly connected to high availability and self-healing: Kubernetes does not “move” Pods from a dead node; it replaces them. The eviction timeout gives the cluster time to confirm the node is truly unavailable, avoiding flapping in unstable networks. Once eviction begins, replacement Pods can be scheduled elsewhere (assuming capacity exists), which is the normal recovery path for stateless workloads.

It’s also worth noting that graceful operational handling can be influenced by PodDisruptionBudgets (for voluntary disruptions) and by workload design (replicas across nodes/zones). But the question is testing the default timer value, which is five minutes in this context.

Therefore, among the choices provided, the correct answer is D.

Question: 188

What is the main purpose of a DaemonSet?

- A. A DaemonSet ensures that all (or certain) nodes run a copy of a Pod.
- B. A DaemonSet ensures that the kubelet is constantly up and running.
- C. A DaemonSet ensures that there are as many pods running as specified in the replicas field.
- D. A DaemonSet ensures that a process (agent) runs on every node.

Answer: A

Explanation:

The correct answer is A. A DaemonSet is a workload controller whose job is to ensure that a specific Pod runs on all nodes (or on a selected subset of nodes) in the cluster. This is fundamentally different from Deployments/ReplicaSets, which aim to maintain a certain replica count regardless of node count. With a DaemonSet, the number of Pods is

implicitly tied to the number of eligible nodes: add a node, and the DaemonSet automatically schedules a Pod there; remove a node, and its Pod goes away.

DaemonSets are commonly used for node-level services and background agents: log collectors, node monitoring agents, storage daemons, CNI components, or security agents—anything where you want a presence on each node to interact with node resources. This aligns with option D’s phrasing (“agent on every node”), but option A is the canonical definition and is slightly broader because it covers “all or certain nodes” (via node selectors/affinity/taints-tolerations) and the fact that the unit is a Pod.

Why the other options are wrong: DaemonSets do not “keep kubelet running” (B); kubelet is a node service managed by the OS. DaemonSets do not use a replicas field to maintain a specific count (C); that’s Deployment/ReplicaSet behavior.

Operationally, DaemonSets matter for cluster operations because they provide consistent node coverage and automatically react to node pool scaling. They also require careful scheduling constraints so they land only where intended (e.g., only Linux nodes, only GPU nodes). But the main purpose remains: ensure a copy of a Pod runs on each relevant node—option A.

Question: 189

Why do administrators need a container orchestration tool?

- A. To manage the lifecycle of an elevated number of containers.
- B. To assess the security risks of the container images used in production.
- C. To learn how to transform monolithic applications into microservices.
- D. Container orchestration tools such as Kubernetes are the future.

Answer: A

Explanation:

The correct answer is A. Container orchestration exists because running containers at scale is hard: you need to schedule workloads onto machines, keep them healthy, scale them up and down, roll out updates safely, and recover from failures automatically. Administrators (and platform teams) use orchestration tools like Kubernetes to manage the lifecycle of many containers across many nodes—handling placement, restart, rescheduling, networking/service discovery, and desired-state reconciliation.

At small scale, you can run containers manually or with basic scripts. But at “elevated” scale (many services, many replicas, many nodes), manual management becomes unreliable and brittle. Orchestration provides primitives and controllers that continuously converge actual state toward desired state: if a container crashes, it is restarted; if a node dies, replacement Pods are scheduled; if traffic increases, replicas can be increased via autoscaling; if configuration changes, rolling updates can be coordinated with readiness checks.

Option B (security risk assessment) is important, but it’s not why orchestration tools exist. Image scanning and supply-chain security are typically handled by CI/CD tooling and registries, not by orchestration as the primary purpose. Option C is a separate architectural modernization effort; orchestration can support microservices, but it isn’t required “to learn transformation.” Option D is an opinion statement rather than a functional need.

So the core administrator need is lifecycle management at scale: ensuring workloads run reliably, predictably, and efficiently across a fleet. That is exactly what option A states.

Question: 190

Which two elements are shared between containers in the same pod?

- A. Network resources and liveness probes.
- B. Storage and container image registry.
- C. Storage and network resources.
- D. Network resources and Dockerfiles.

Answer: C

Explanation:

The correct answer is C: Storage and network resources. In Kubernetes, a Pod is the smallest schedulable unit and acts like a “logical host” for its containers. Containers inside the same Pod share a number of namespaces and resources, most notably:

Network: all containers in a Pod share the same network namespace, which means they share a single Pod IP address and the same port space. They can talk to each other via localhost and coordinate tightly without exposing separate network endpoints.

Storage: containers in a Pod can share data through Pod volumes. Volumes (like emptyDir, ConfigMap/Secret volumes, or PVC-backed volumes) are defined at the Pod level and can be mounted into multiple containers within the Pod. This enables common patterns like a sidecar writing logs to a shared volume that the main container generates, or an init/sidecar container producing configuration or certificates for the main container.

Why other options are wrong: liveness probes (A) are defined per container (or per Pod template) but are not a “shared” resource between containers. A container image registry (B) is an external system and not a shared in-Pod element. Dockerfiles (D) are build-time artifacts, irrelevant at runtime, and not shared resources.

This question is a classic test of Pod fundamentals: multi-container Pods work precisely because they share networking and volumes. This is also why the sidecar pattern is feasible—sidecars can intercept traffic on localhost, export metrics, or ship logs while sharing the same lifecycle boundary and scheduling placement.

Therefore, the verified correct choice is C.

Question: 191

In Kubernetes, if the API version of feature is v2beta3, it means that:

- A. The version will remain available for all future releases within a Kubernetes major version.
- B. The API may change in incompatible ways in a later software release without notice.
- C. The software is well tested. Enabling a feature is considered safe.
- D. The software may contain bugs. Enabling a feature may expose bugs.

Answer: B

Explanation:

The correct answer is B. In Kubernetes API versioning, the stability level is encoded in the version string: alpha, beta, and stable (v1). A version like v2beta3 indicates the API is in a beta stage. Beta APIs are more mature than alpha, but they are not fully guaranteed stable in perpetuity the way v1 stable APIs are intended to be. The key implication is that while beta APIs are generally usable, they can still undergo incompatible changes in future releases as the API design evolves.

Option B captures that meaning: a beta API may change in ways that break compatibility. This is why teams should treat beta APIs with some caution in production: verify upgrade plans, monitor deprecation notices, and be prepared to adjust manifests or client code when moving between Kubernetes versions.

Why the other options are incorrect:

A implies permanence across all future releases in a major version, which is not a beta guarantee. Kubernetes has deprecation and graduation processes, but beta does not equal “forever.”

C overstates safety; beta is typically “tested and enabled by default” for some features, but it’s not the same as stable API guarantees.

D is too vague and misaligned. While any software may contain bugs, the defining point of “beta API” is about stability/compatibility guarantees, not merely “bugs.”

In practice, Kubernetes communicates API lifecycle clearly: alpha is experimental and may be disabled by default; beta is feature-complete-ish but may change; stable v1 is strongly compatibility- focused with formal deprecation policies. So, a v2beta3 API signals: usable, but not fully locked— hence B.

Question: 192

What is the API that exposes resource metrics from the metrics-server?

- A. custom.k8s.io
- B. resources.k8s.io
- C. metrics.k8s.io
- D. cadvisor.k8s.io

Answer: C

Explanation:

The correct answer is C: metrics.k8s.io. Kubernetes’ metrics-server is the standard component that provides resource metrics (primarily CPU and memory) for nodes and pods. It aggregates this information (sourced from kubelet/cAdvisor) and serves it through the Kubernetes aggregated API under the group metrics.k8s.io. This is what enables commands like `kubectl top nodes` and `kubectl top pods`, and it is also a key data source for autoscaling with the Horizontal Pod Autoscaler (HPA) when scaling on CPU/memory utilization.

Why the other options are wrong:

custom.k8s.io is not the standard API group for metrics-server resource metrics. Custom metrics are typically served through the custom metrics API (commonly custom.metrics.k8s.io) via adapters (e.g., Prometheus Adapter), not metrics-server.

resources.k8s.io is not the metrics-server API group.

cadvisor.k8s.io is not exposed as a Kubernetes aggregated metrics API. cAdvisor is a component integrated into kubelet that provides container stats, but metrics-server is the thing that exposes the aggregated Kubernetes metrics API, and the canonical group is metrics.k8s.io.

Operationally, it's important to understand the boundary: metrics-server provides basic resource metrics suitable for core autoscaling and "top" views, but it is not a full observability system (it does not store long-term metrics history like Prometheus). For richer metrics (SLOs, application metrics, long-term trending), teams typically deploy Prometheus or a managed monitoring backend. Still, when the question asks specifically which API exposes metrics-server data, the answer is definitively metrics.k8s.io.

Question: 193

Which of the following resources helps in managing a stateless application workload on a Kubernetes cluster?

- A. DaemonSet
- B. StatefulSet
- C. kubectI
- D. Deployment

Answer: D

Explanation:

The correct answer is D: Deployment. A Deployment is the standard Kubernetes controller for managing stateless applications. It provides declarative updates, replica management, and rollout/rollback functionality. You define the

desired state (container image, environment variables, ports, replica count) in the Deployment spec, and Kubernetes ensures the specified number of Pods are running and updated according to strategy (RollingUpdate by default).

Stateless workloads are ideal for Deployments because each replica is interchangeable. If a Pod dies, a new one can be created anywhere; if traffic increases, replicas can be increased; if you need to update the app, a new ReplicaSet is created and traffic shifts gradually to new Pods. Deployments integrate naturally with Services for stable networking and load balancing.

Why the other options are incorrect:

A DaemonSet ensures one Pod per node (or selected nodes). It's for node-level agents, not generic stateless service replicas.

A StatefulSet is for workloads needing stable identity, ordered rollout, and persistent storage per replica (databases, quorum systems). That's not the typical stateless app case.

kubectl is a CLI tool; it doesn't "manage" workloads as a controller resource.

In real cluster operations, almost every stateless microservice is represented as a Deployment plus a Service (and often an Ingress/Gateway for edge routing). Deployments also support advanced delivery patterns (maxSurge/maxUnavailable tuning) and easy integration with HPA for horizontal scaling. Because the question is specifically "managing a stateless application workload," the Kubernetes resource designed for that is clearly the Deployment.

Question: 194

Which mechanism can be used to automatically adjust the amount of resources for an application?

- A. Horizontal Pod Autoscaler (HPA)
- B. Kubernetes Event-driven Autoscaling (KEDA)
- C. Cluster Autoscaler
- D. Vertical Pod Autoscaler (VPA)

Answer: A

Explanation:

The verified answer in the PDF is A (HPA), and that aligns with the common Kubernetes meaning of “adjust resources for an application” by scaling replicas. The Horizontal Pod Autoscaler automatically

changes the number of Pod replicas for a workload (typically a Deployment) based on observed metrics such as CPU utilization, memory (in some configurations), or custom/external metrics. By increasing replicas under load, the application gains more total CPU/memory capacity available across Pods; by decreasing replicas when load drops, it reduces resource consumption and cost.

It’s important to distinguish what each mechanism adjusts:

HPA adjusts replica count (horizontal scaling).

VPA adjusts Pod resource requests/limits (vertical scaling), which is literally “amount of CPU/memory per pod,” but it often requires restarts to apply changes depending on mode.

Cluster Autoscaler adjusts the number of nodes in the cluster, not application replicas.

KEDA is event-driven autoscaling that often drives HPA behavior using external event sources (queues, streams), but it’s not the primary built-in mechanism referenced in many foundational Kubernetes questions.

Given the wording and the provided answer key, the intended interpretation is: “automatically adjust the resources available to the application” by scaling out/in the number of replicas. That’s exactly HPA’s role. For example, if CPU utilization exceeds a target (say 60%), HPA computes a higher desired replica count and updates the workload. The Deployment then creates more Pods, distributing load and increasing available compute.

So, within this question set, the verified correct choice is A (Horizontal Pod Autoscaler).

Question: 195

Which of the following is a recommended security habit in Kubernetes?

- A. Run the containers as the user with group ID 0 (root) and any user ID.
- B. Disallow privilege escalation from within a container as the default option.
- C. Run the containers as the user with user ID 0 (root) and any group ID.
- D. Allow privilege escalation from within a container as the default option.

Answer: B

Explanation:

The correct answer is B. A widely recommended Kubernetes security best practice is to disallow privilege escalation inside containers by default. In Kubernetes Pod/Container security context, this is represented by `allowPrivilegeEscalation: false`. This setting prevents a process from gaining more privileges than its parent process—commonly via `setuid/setgid` binaries or other privilege-escalation mechanisms. Disallowing privilege escalation reduces the blast radius of a compromised container and aligns with least-privilege principles.

Options A and C are explicitly unsafe because they encourage running as root (UID 0 and/or GID 0). Running containers as root increases risk: if an attacker breaks out of the application process or exploits kernel/runtime vulnerabilities, having root inside the container can make privilege escalation and lateral movement easier. Modern Kubernetes security guidance strongly favors running as non-root (`runAsNonRoot: true`, explicit `runAsUser`), dropping Linux capabilities, using read-only root filesystems, and applying restrictive `seccomp/AppArmor/SELinux` profiles where possible.

Option D is the opposite of best practice. Allowing privilege escalation by default increases the attack surface and violates the idea of secure defaults.

Operationally, this habit is often enforced via admission controls and policies (e.g., Pod Security Admission in “restricted” mode, or policy engines like OPA Gatekeeper/Kyverno). It’s also important for compliance: many security baselines require containers to run as non-root and to prevent privilege escalation.

So, the recommended security habit among the choices is clearly B: Disallow privilege escalation.

Question: 196

What are the 3 pillars of Observability?

- A. Metrics, Logs, and Traces
- B. Metrics, Logs, and Spans
- C. Metrics, Data, and Traces
- D. Resources, Logs, and Tracing

Answer: A

Explanation:

The correct answer is A: Metrics, Logs, and Traces. These are widely recognized as the “three pillars” because together they provide complementary views into system behavior:

Metrics are numeric time series collected over time (CPU usage, request rate, error rate, latency percentiles). They are best for dashboards, alerting, and capacity planning because they are structured and aggregatable. In Kubernetes, metrics underpin autoscaling and operational visibility (node/pod resource usage, cluster health signals).

Logs are discrete event records (often text) emitted by applications and infrastructure components. Logs provide detailed context for debugging: error messages, stack traces, warnings, and business events. In Kubernetes, logs are commonly collected from container stdout/stderr and aggregated centrally for search and correlation.

Traces capture the end-to-end journey of a request through a distributed system, breaking it into spans. Tracing is crucial in microservices because a single user request may cross many services; traces show where latency accumulates and which dependency fails. Tracing also enables root cause analysis when metrics indicate degradation but don't pinpoint the culprit.

Why the other options are wrong: a span is a component within tracing, not a top-level pillar; “data” is too generic; and “resources” are not an observability signal category. The pillars are defined by signal type and how they're used operationally.

In cloud-native practice, these pillars are often unified via correlation IDs and shared context: metrics alerts link to logs and traces for the same timeframe/request. Tooling like Prometheus (metrics), log aggregators (e.g., Loki/Elastic), and tracing systems (Jaeger/Tempo/OpenTelemetry) work together to provide a complete observability story.

Therefore, the verified correct answer is A.

Question: 197

What edge and service proxy tool is designed to be integrated with cloud native applications?

- A. CoreDNS
- B. CNI
- C. gRPC
- D. Envoy

Answer: D

Explanation:

The correct answer is D: Envoy. Envoy is a high-performance edge and service proxy designed for cloud-native environments. It is commonly used as the data plane in service meshes and modern API gateways because it provides consistent traffic management, observability, and security features across microservices without requiring every application to implement those capabilities directly.

Envoy operates at Layer 7 (application-aware) and supports protocols like HTTP/1.1, HTTP/2, gRPC, and more. It can handle routing, load balancing, retries, timeouts, circuit breaking, rate limiting, TLS termination, and mutual TLS (mTLS). Envoy also emits rich telemetry (metrics, access logs, tracing) that integrates well with cloud-native observability stacks.

Why the other options are incorrect:

CoreDNS (A) provides DNS-based service discovery within Kubernetes; it is not an edge/service proxy.

CNI (B) is a specification and plugin ecosystem for container networking (Pod networking), not a proxy.

gRPC (C) is an RPC protocol/framework used by applications; it's not a proxy tool. (Envoy can proxy gRPC traffic, but gRPC itself isn't the proxy.)

In Kubernetes architectures, Envoy often appears in two places: (1) at the edge as part of an ingress/gateway layer, and (2) sidecar proxies alongside Pods in a service mesh (like Istio) to standardize service-to-service communication controls and telemetry. This is why it is described as "designed to be integrated with cloud native applications": it's purpose-built for dynamic service discovery, resilient routing, and operational visibility in distributed systems.

So the verified correct choice is D (Envoy).

Question: 198

What is Flux constructed with?

- A. GitLab Environment Toolkit
- B. GitOps Toolkit
- C. Helm Toolkit

D. GitHub Actions Toolkit

Answer: B

Explanation:

The correct answer is B: GitOps Toolkit. Flux is a GitOps solution for Kubernetes, and in Flux v2 the project is built as a set of Kubernetes controllers and supporting components collectively referred to as the GitOps Toolkit. This toolkit provides the building blocks for implementing GitOps reconciliation: sourcing artifacts (Git repositories, Helm repositories, OCI artifacts), applying manifests (Kustomize/Helm), and continuously reconciling cluster state to match the desired state declared in Git.

This construction matters because it reflects Flux's modular architecture. Instead of being a single monolithic daemon, Flux is composed of controllers that each handle a part of the GitOps workflow: fetching sources, rendering configuration, and applying changes. This makes it more Kubernetes- native: everything is declarative, runs in the cluster, and can be managed like other workloads (RBAC, namespaces, upgrades, observability).

Why the other options are wrong:

"GitLab Environment Toolkit" and "GitHub Actions Toolkit" are not what Flux is built from. Flux can integrate with many SCM providers and CI systems, but it is not "constructed with" those.

"Helm Toolkit" is not the named foundational set Flux is built upon. Flux can deploy Helm charts, but that's a capability, not its underlying construction.

In cloud-native delivery, Flux implements the key GitOps control loop: detect changes in Git (or other declared sources), compute desired Kubernetes state, and apply it while continuously checking for drift. The GitOps Toolkit is the set of controllers enabling that loop.

Therefore, the verified correct answer is B.

Question: 199

In Kubernetes, which abstraction defines a logical set of Pods and a policy by which to access them?

A. Service Account

- B. NetworkPolicy
- C. Service
- D. Custom Resource Definition

Answer: C

Explanation:

The correct answer is C: Service. A Kubernetes Service is an abstraction that provides stable access to a logical set of Pods. Pods are ephemeral: they can be rescheduled, recreated, and scaled, which changes their IP addresses over time. A Service solves this by providing a stable identity—typically a virtual IP (ClusterIP) and a DNS name—and a traffic-routing policy that directs requests to the current set of backend Pods.

Services commonly select Pods using labels via a selector (e.g., `app=web`). Kubernetes then maintains the backend endpoint list (Endpoints/EndpointSlices). The cluster networking layer routes traffic sent to the Service IP/port to one of the Pod endpoints, enabling load distribution across replicas. This is fundamental to microservices architectures: clients call the Service name, not individual Pods.

Why the other options are incorrect:

A ServiceAccount is an identity for Pods to authenticate to the Kubernetes API; it doesn't define a set of Pods nor traffic access policy.

A NetworkPolicy defines allowed network flows (who can talk to whom) but does not provide stable addressing or load-balanced access to Pods. It is a security policy, not an exposure abstraction.

A CustomResourceDefinition extends the Kubernetes API with new resource types; it's unrelated to service discovery and traffic routing for a set of Pods.

Understanding Services is core Kubernetes fundamentals: they decouple backend Pod churn from client connectivity. Services also integrate with different exposure patterns via type (ClusterIP, NodePort, LoadBalancer, ExternalName) and can be paired with Ingress/Gateway for HTTP routing. But the essential definition in the question—"logical set of Pods and a policy to access them"—is exactly the textbook description of a Service.

Therefore, the verified correct answer is C.

Question: 200

Which field in a Pod or Deployment manifest ensures that Pods are scheduled only on nodes with specific labels?

A.

resources: disktype: ssd B.

labels: disktype: ssd C.

nodeSelector: disktype: ssd D.

annotations:

disktype: ssd

Answer: C

Explanation:

In Kubernetes, Pod scheduling is handled by the Kubernetes scheduler, which is responsible for assigning Pods to suitable nodes based on a set of constraints and policies. One of the simplest and most commonly used mechanisms to control where Pods are scheduled is the nodeSelector field. The nodeSelector field allows you to constrain a Pod so that it is only eligible to run on nodes that have specific labels.

Node labels are key–value pairs attached to nodes by cluster administrators or automation tools. These labels typically describe node characteristics such as hardware type, disk type, geographic zone, or environment. For example, a node might be labeled with disktype=ssd to indicate that it has SSD-backed storage. When a Pod specification includes a nodeSelector with the same key–value pair, the scheduler will only consider nodes that match this label when placing the Pod.

Option A (resources) is incorrect because resource specifications are used to define CPU and memory requests and limits for containers, not to influence node selection based on labels. Option B (labels) is also incorrect because Pod labels are metadata used for identification, grouping, and selection by other Kubernetes objects such as Services and Deployments; they do not affect scheduling decisions. Option D (annotations) is incorrect because annotations are intended for storing non-identifying metadata and are not interpreted by the scheduler for placement decisions.

The nodeSelector field is evaluated during scheduling, and if no nodes match the specified labels, the Pod will remain in a Pending state. While nodeSelector is simple and effective, it is considered a basic scheduling mechanism. For more advanced scheduling requirements—such as expressing preferences, using set-based matching, or combining multiple conditions—Kubernetes also provides node affinity and anti-affinity. However, nodeSelector remains a foundational and widely used feature for enforcing strict node placement based on labels, making option C the correct and verified answer according to Kubernetes documentation.

Question: 201

Which of the following actions is supported when working with Pods in Kubernetes?

- A. Managing static Pods directly through the API server.
- B. Guaranteeing Pods always stay on the same node once scheduled.
- C. Renaming containers in a Pod using kubectl patch.
- D. Creating Pods through workload resources like Deployments.

Answer: D

Explanation:

In Kubernetes, Pods are the smallest deployable units and represent one or more containers that share networking and storage. While Pods can be created directly, Kubernetes strongly encourages users to manage Pods indirectly through higher-level workload resources. Among the options provided, creating Pods through workload resources like Deployments is a fully supported and recommended practice.

Workload resources such as Deployments, ReplicaSets, StatefulSets, and Jobs are designed to manage Pods declaratively. A Deployment, for example, defines a desired state—such as the number of replicas and the Pod template—and Kubernetes continuously works to maintain that state. If a Pod crashes, is deleted, or a node fails, the Deployment automatically creates a replacement Pod. This model provides self-healing, scalability, rolling updates, and rollback capabilities, which are not available when managing standalone Pods.

Option A is incorrect because static Pods are not managed through the API server. Static Pods are created and managed directly by the kubelet on a specific node using manifest files placed on disk. Although the API server becomes aware of static Pods, they cannot be created, modified, or deleted through it.

Option B is incorrect because Kubernetes does not guarantee that Pods will always remain on the same node. If a node becomes unhealthy or a Pod is evicted, the scheduler may place a replacement Pod on a different node. Only certain workload patterns, such as StatefulSets with persistent storage, attempt to preserve identity—not node placement.

Option C is also incorrect because container names within a Pod are immutable. Kubernetes does not allow renaming containers using `kubectl patch` or any other mechanism after the Pod has been created.

Therefore, the correct and verified answer is option D: creating Pods through workload resources like Deployments, which aligns with Kubernetes design principles and official documentation.

Question: 202

A site reliability engineer needs to temporarily prevent new Pods from being scheduled on node-2 while keeping the existing workloads running without disruption. Which `kubectl` command should be used?

- A. `kubectl cordon node-2`
- B. `kubectl delete node-2`
- C. `kubectl drain node-2`
- D. `kubectl pause deployment`

Answer: A

Explanation:

In Kubernetes, node maintenance and availability are common operational tasks, and the platform provides specific commands to control how the scheduler places Pods on nodes. When the requirement is to temporarily prevent new Pods from being scheduled on a node without affecting the currently running Pods, the correct approach is to cordon the node.

The command `kubectl cordon node-2` marks the node as unschedulable. This means the Kubernetes scheduler will no longer place any new Pods onto that node. Importantly, cordoning a node does not evict, restart, or interrupt existing Pods. All workloads already running on the node continue operating normally. This makes cordoning ideal for scenarios such as diagnostics, monitoring, or preparing for future maintenance while ensuring zero workload disruption.

Option B, `kubectl delete node-2`, is incorrect because deleting a node removes it entirely from the cluster. This action would cause Pods running on that node to be terminated and rescheduled elsewhere, resulting in disruption—exactly what the question specifies must be avoided.

Option C, `kubectl drain node-2`, is also incorrect in this context. Draining a node safely evicts Pods (except for certain exclusions like DaemonSets) and reschedules them onto other nodes. While drain is useful for maintenance and upgrades, it does not keep existing workloads running on the node, making it unsuitable here.

Option D, `kubectl pause deployment`, applies only to Deployments and merely pauses rollout updates. It does not affect node-level scheduling behavior and has no impact on where Pods are placed by the scheduler.

Therefore, the correct and verified answer is Option A: `kubectl cordon node-2`, which aligns with Kubernetes operational best practices and official documentation for non-disruptive node management.

Question: 203

What is an advantage of using the Gateway API compared to Ingress in Kubernetes?

- A. To automatically scale workloads based on CPU and memory utilization.
- B. To provide clearer role separation between infrastructure providers and application developers.
- C. To configure routing rules through annotations directly on Ingress resources.
- D. To expose an application externally by creating only a Service resource.

Answer: B

Explanation:

The Gateway API is a newer Kubernetes networking API designed to address several limitations of the traditional Ingress resource. One of its most significant advantages is the clear separation of roles and responsibilities between infrastructure providers (such as platform teams or cluster administrators) and application developers. This design principle is a core motivation behind the Gateway API and directly differentiates it from Ingress.

With Ingress, a single resource often combines concerns such as load balancer configuration, TLS settings, routing rules, and application-level details. This frequently leads to heavy reliance on annotations, which are controller-specific, non-standardized, and blur ownership boundaries. Application developers may need elevated permissions to modify Ingress objects, even when changes affect shared infrastructure, creating operational risk.

The Gateway API introduces multiple distinct resources—such as GatewayClass, Gateway, and route resources (e.g., HTTPRoute)—each aligned with a specific role. Infrastructure providers manage GatewayClass and Gateway resources, which define how traffic enters the cluster and what capabilities are available. Application developers interact primarily with route resources to define how traffic is routed to their Services, without needing access to the underlying infrastructure

configuration. This separation improves security, governance, and scalability in multi-team environments.

Option A is incorrect because automatic scaling based on CPU and memory is handled by the Horizontal Pod Autoscaler, not by Gateway API or Ingress. Option C describes a characteristic of Ingress, not an advantage of Gateway API; in fact, Gateway API explicitly reduces reliance on annotations by using structured, portable fields. Option D is incorrect because exposing applications externally requires more than just a Service; traffic management resources like Ingress or Gateway are still necessary.

Therefore, the correct and verified answer is Option B, as the Gateway API's role-oriented design is a key advancement over Ingress and is clearly documented in Kubernetes networking architecture guidance.

Question: 204

A Pod is stuck in the CrashLoopBackOff state. Which is the correct way to troubleshoot this issue?

- A. Use `kubectl exec <pod-name> -- bash` to connect inside the container and then check system logs in `/var/log/kubelet.log`.
- B. Use `kubectl describe pod <pod-name>` to review recent events and then `kubectl logs <pod-name>` to inspect container output.

- C. Use `kubectl get nodes` to verify node capacity and then `kubectl apply -f <pod.yaml>` to restart the Pod.
- D. Use `kubectl top pod <pod-name>` to check CPU usage and then scale the Deployment to more replicas.

Answer: B

Explanation:

The `CrashLoopBackOff` state in Kubernetes indicates that a container inside a Pod is repeatedly starting, crashing, and then being restarted by the kubelet with increasing backoff delays. This is typically caused by application-level issues such as misconfiguration, missing environment variables,

failed startup commands, application crashes, or incorrect container images. Proper troubleshooting focuses on identifying why the container is failing shortly after startup.

The most effective and recommended approach is to first use `kubectl describe pod <pod-name>`. This command provides detailed information about the Pod, including its current state, restart count, container statuses, and—most importantly—the Events section. Events often reveal critical clues such as image pull errors, failed health checks, permission issues, or failed command executions. These messages are generated by Kubernetes components and are essential for understanding the failure context.

After reviewing the events, the next step is to inspect the container's logs using `kubectl logs <podname>`. Container logs typically capture application output written to standard output and standard error. For a crashing container, these logs often show stack traces, configuration errors, or explicit failure messages that explain why the process exited. If the container restarts too quickly, logs from the previous run can be retrieved using the `--previous` flag.

Option A is incorrect because `kubectl exec` usually fails when containers are repeatedly crashing, and `/var/log/kubelet.log` is a node-level log not accessible from inside the container. Option C is incorrect because reapplying the Pod manifest does not address the underlying crash cause. Option D focuses on resource usage and scaling, which does not resolve application startup failures.

Therefore, the correct and verified answer is Option B, which aligns with Kubernetes documentation and best practices for diagnosing `CrashLoopBackOff` conditions.

Question: 205

Ceph is a highly scalable distributed storage solution for block storage, object storage, and shared filesystems with years of production deployments. Which open-source cloud native storage orchestrator automates deployment and management of Ceph to provide self-managing, selfscaling, and self-healing storage services?

A. CubeFS

B. OpenEBS

C. Rook

D. MinIO

Answer: C

Explanation:

Rook is the open-source, cloud-native storage orchestrator specifically designed to automate the deployment, configuration, and lifecycle management of Ceph within Kubernetes environments. Its primary goal is to transform complex, traditionally manual storage systems like Ceph into Kubernetes-native services that are easy to operate and highly resilient.

Ceph itself is a mature and powerful distributed storage platform that supports block storage (RBD), object storage (RGW), and shared filesystems (CephFS). However, operating Ceph directly requires deep expertise, careful configuration, and continuous operational management. Rook addresses this challenge by running Ceph as a set of Kubernetes-managed components and exposing storage capabilities through Kubernetes Custom Resource Definitions (CRDs). This allows administrators to declaratively define storage clusters, pools, filesystems, and object stores using familiar Kubernetes patterns.

Rook continuously monitors the health of the Ceph cluster and takes automated actions to maintain the desired state. If a Ceph daemon fails or a node becomes unavailable, Rook works with Kubernetes scheduling and Ceph's internal replication mechanisms to ensure data durability and service continuity. This enables self-healing behavior. Scaling storage capacity is also simplified—adding nodes or disks allows Rook and Ceph to automatically rebalance data, providing self-scaling capabilities without manual intervention.

The other options are incorrect for this use case. CubeFS is a distributed filesystem but is not a Ceph orchestrator.

OpenEBS focuses on container-attached storage and local or replicated volumes rather than managing Ceph itself. MinIO is an object storage server compatible with S3 APIs, but it does not orchestrate Ceph or provide block and filesystem services.

Therefore, the correct and verified answer is Option C: Rook, which is the officially recognized Kubernetes-native orchestrator for Ceph, delivering automated, resilient, and scalable storage management aligned with cloud-native principles.

Question: 206

A platform engineer wants to ensure that a new microservice is automatically deployed to every cluster registered in Argo CD. Which configuration best achieves this goal?

- A. Set up a Kubernetes CronJob that redeploys the microservice to all registered clusters on a schedule.
- B. Manually configure every registered cluster with the deployment YAML for installing the microservice.
- C. Create an Argo CD ApplicationSet that uses a Git repository containing the microservice manifests.
- D. Use a Helm chart to package the microservice and manage it with a single Application defined in Argo CD.

Answer: C

Explanation:

Argo CD is a declarative GitOps continuous delivery tool designed to manage Kubernetes applications across one or many clusters. When the requirement is to automatically deploy a microservice to every cluster registered in Argo CD, the most appropriate and scalable solution is to use an ApplicationSet.

The ApplicationSet controller extends Argo CD by enabling the dynamic generation of multiple Argo CD Applications from a single template. One of its most powerful features is the cluster generator, which automatically discovers all clusters registered with Argo CD and creates an Application for each of them. By combining this generator with a Git repository containing the microservice manifests, the platform engineer ensures that the microservice is consistently deployed to all existing clusters—and any new clusters added in the future—without manual intervention.

This approach aligns perfectly with GitOps principles. The desired state of the microservice is defined once in Git, and Argo CD continuously reconciles that state across all target clusters. Any updates to the microservice manifests are automatically rolled out everywhere in a controlled and auditable manner. This provides strong guarantees around consistency, scalability, and operational simplicity.

Option A is incorrect because a CronJob introduces imperative redeployment logic and does not integrate with Argo CD's reconciliation model. Option B is not scalable or maintainable, as it requires manual configuration for each cluster and increases the risk of configuration drift. Option D, while useful for packaging applications, still results in a single Application object and does not natively handle multi-cluster fan-out by itself.

Therefore, the correct and verified answer is Option C: creating an Argo CD ApplicationSet backed by a Git repository,

which is the recommended and documented solution for multi-cluster application delivery in Argo CD.

Question: 207

What is the Kubernetes abstraction that allows groups of Pods to be exposed inside a Kubernetes cluster?

- A. Deployment
- B. Daemon
- C. Unit
- D. Service

Answer: D

Explanation:

In Kubernetes, Pods are ephemeral by design. They can be created, destroyed, rescheduled, or replaced at any time, and each Pod receives its own IP address. Because of this dynamic nature, directly relying on Pod IPs for communication is unreliable. To solve this problem, Kubernetes provides the Service abstraction, which allows a stable way to expose and access a group of Pods inside (and sometimes outside) the cluster.

A Service defines a logical set of Pods using label selectors and provides a consistent virtual IP address and DNS name for accessing them. Even if individual Pods fail or are replaced, the Service remains stable, and traffic is automatically routed to healthy Pods that match the selector. This makes Services a fundamental building block for internal communication between applications within a Kubernetes cluster.

Deployments (Option A) are responsible for managing the lifecycle of Pods, including scaling, rolling updates, and self-healing. However, Deployments do not provide networking or exposure capabilities. They control how Pods run, not how they are accessed.

Option B, "Daemon," is not a valid Kubernetes resource. The correct resource is a DaemonSet, which ensures that a copy of a Pod runs on each (or selected) node in the cluster. DaemonSets are used for node-level workloads like logging or monitoring agents, not for exposing Pods.

Option C, "Unit," is not a Kubernetes concept at all and does not exist in Kubernetes architecture.

Services can be configured in different ways depending on access requirements, such as ClusterIP for internal access,

NodePort or LoadBalancer for external access, and Headless Services for direct Pod discovery. Regardless of type, the core purpose of a Service is to expose a group of Pods in a stable and reliable way.

Therefore, the correct and verified answer is Option D: Service, which is the Kubernetes abstraction specifically designed to expose groups of Pods within a cluster.

Question: 208

What are the two essential operations that the kube-scheduler normally performs?

- A. Pod eviction or starting
- B. Resource monitoring and reporting
- C. Filtering and scoring nodes
- D. Starting and terminating containers

Answer: C

Explanation:

The kube-scheduler is a core control plane component in Kubernetes responsible for assigning newly created Pods to appropriate nodes. Its primary responsibility is decision-making, not execution. To make an informed scheduling decision, the kube-scheduler performs two essential operations: **filtering and scoring nodes**.

The scheduling process begins when a Pod is created without a node assignment. The scheduler first evaluates all available nodes and applies a set of filtering rules. During this phase, nodes that do not meet the Pod's requirements are eliminated. Filtering criteria include resource availability (CPU and memory requests), node selectors, node affinity rules, taints and tolerations, volume constraints, and other policy-based conditions. Any node that fails one or more of these checks is excluded from consideration.

Once filtering is complete, the scheduler moves on to the scoring phase. In this step, each remaining eligible node is assigned a score based on a collection of scoring plugins. These plugins evaluate

factors such as resource utilization balance, affinity preferences, topology spread constraints, and custom scheduling policies. The purpose of scoring is to rank nodes according to how well they satisfy the Pod's placement preferences. The node with the highest total score is selected as the best candidate.

Option A is incorrect because Pod eviction is handled by other components such as the kubelet and controllers, and starting Pods is the responsibility of the kubelet. Option B is incorrect because resource monitoring and reporting are performed by components like metrics-server, not the scheduler. Option D is also incorrect because starting and terminating containers is entirely handled by the kubelet and the container runtime.

By separating filtering (eligibility) from scoring (preference), the kube-scheduler provides a flexible, extensible, and policy-driven scheduling mechanism. This design allows Kubernetes to support diverse workloads and advanced placement strategies while maintaining predictable scheduling behavior.

Therefore, the correct and verified answer is Option C: Filtering and scoring nodes, as documented in Kubernetes scheduling architecture.

Question: 209

In Kubernetes, what is the primary responsibility of the kubelet running on each worker node?

- A. To allocate persistent storage volumes and manage distributed data replication for Pods.
- B. To manage cluster state information and handle all scheduling decisions for workloads.
- C. To ensure that containers defined in Pod specifications are running and remain healthy on the node.
- D. To provide internal DNS resolution and route service traffic between Pods and nodes.

Answer: C

Explanation:

The kubelet is the primary node-level agent in Kubernetes and plays a critical role in ensuring that workloads run correctly on each worker node. Its main responsibility is to ensure that the containers

described in Pod specifications are running and remain healthy on that node, which makes option C the correct answer.

Once the Kubernetes scheduler assigns a Pod to a node, the kubelet on that node takes over execution responsibilities. It watches the API server for Pod specifications that are scheduled to its node and then interacts with the container runtime to start, stop, and manage the containers defined in those Pods. The kubelet continuously monitors container health and reports Pod and node status back to the API server, enabling Kubernetes to make informed decisions about

restarts, rescheduling, or remediation.

Health checks are another key responsibility of the kubelet. It executes liveness, readiness, and startup probes as defined in the Pod specification. Based on probe results, the kubelet may restart containers or update Pod status to reflect whether the application is ready to receive traffic. This behavior directly supports Kubernetes' self-healing capabilities.

Option A is incorrect because persistent storage allocation and data replication are handled by storage systems, CSI drivers, and controllers—not by the kubelet itself. Option B is incorrect because cluster state management and scheduling decisions are the responsibility of control plane components such as the API server, controller manager, and kube-scheduler. Option D is incorrect because DNS resolution and service traffic routing are handled by components like CoreDNS and kube-proxy.

In summary, the kubelet acts as the “node supervisor” for Kubernetes workloads. By ensuring containers are running as specified and continuously reporting their status, the kubelet forms the essential link between the Kubernetes control plane and the actual execution of applications on worker nodes. This clearly aligns with Option C as the correct and verified answer.

Question: 210

In Kubernetes, what is the primary responsibility of the kubelet running on each worker node?

- A. To allocate persistent storage volumes and manage distributed data replication for Pods.
- B. To manage cluster state information and handle all scheduling decisions for workloads.
- C. To ensure that containers defined in Pod specifications are running and remain healthy on the node.
- D. To provide internal DNS resolution and route service traffic between Pods and nodes.

Answer: C

Explanation:

The kubelet is a critical Kubernetes component that runs on every worker node and acts as the primary execution agent for Pods. Its core responsibility is to ensure that the containers defined in Pod specifications are running and remain healthy on the node, making option C the correct answer.

Once the Kubernetes scheduler assigns a Pod to a specific node, the kubelet on that node becomes responsible for carrying out the desired state described in the Pod specification. It continuously watches the API server for Pods assigned

to its node and communicates with the container runtime (such as containerd or CRI-O) to start, stop, and restart containers as needed. The kubelet does not make scheduling decisions; it simply executes them.

Health management is another key responsibility of the kubelet. It runs liveness, readiness, and startup probes as defined in the Pod specification. If a container fails a liveness probe, the kubelet restarts it. If a readiness probe fails, the kubelet marks the Pod as not ready, preventing traffic from being routed to it. The kubelet also reports detailed Pod and node status information back to the API server, enabling controllers to take corrective actions when necessary.

Option A is incorrect because persistent volume provisioning and data replication are handled by storage systems, CSI drivers, and controllers—not by the kubelet. Option B is incorrect because cluster state management and scheduling are responsibilities of control plane components such as the API server, controller manager, and kube-scheduler. Option D is incorrect because DNS resolution and service traffic routing are handled by components like CoreDNS and kube-proxy.

In summary, the kubelet serves as the node-level guardian of Kubernetes workloads. By ensuring containers are running exactly as specified and continuously reporting their health and status, the kubelet forms the essential bridge between Kubernetes' declarative control plane and the actual execution of applications on worker nodes.

Question: 211

Which of the following cloud native proxies is used for ingress/egress in a service mesh and can also serve as an application gateway?

- A. Frontend proxy
- B. Kube-proxy
- C. Envoy proxy
- D. Reverse proxy

Answer: C

Explanation:

Envoy Proxy is a high-performance, cloud-native proxy widely used for ingress and egress traffic management in service mesh architectures, and it can also function as an application gateway. It is the foundational data-plane component for popular service meshes such as Istio, Consul, and AWS App Mesh, making option C the correct answer.

In a service mesh, Envoy is typically deployed as a sidecar proxy alongside each application Pod. This allows Envoy to

transparently intercept and manage all inbound and outbound traffic for the service. Through this model, Envoy enables advanced traffic management features such as load balancing, retries, timeouts, circuit breaking, mutual TLS, and fine-grained observability without requiring application code changes.

Envoy is also commonly used at the mesh boundary to handle ingress and egress traffic. When deployed as an ingress gateway, Envoy acts as the entry point for external traffic into the mesh, performing TLS termination, routing, authentication, and policy enforcement. As an egress gateway, it controls outbound traffic from the mesh to external services, enabling security controls and traffic visibility. These capabilities allow Envoy to serve effectively as an application gateway, not just an internal proxy.

Option A, "Frontend proxy," is a generic term and not a specific cloud-native component. Option B, kube-proxy, is responsible for implementing Kubernetes Service networking rules at the node level and does not provide service mesh features or gateway functionality. Option D, "Reverse proxy," is a general architectural pattern rather than a specific cloud-native proxy implementation.

Envoy's extensibility, performance, and deep integration with Kubernetes and service mesh control planes make it the industry-standard proxy for modern cloud-native networking. Its ability to function both as a sidecar proxy and as a centralized ingress or egress gateway clearly establishes Envoy proxy as the correct and verified answer.

Question: 212

What happens if only a limit is specified for a resource and no admission-time mechanism has applied a default request?

- A. Kubernetes will create the container but it will fail with CrashLoopBackOff.
- B. Kubernetes does not allow containers to be created without request values, causing eviction.
- C. Kubernetes copies the specified limit and uses it as the requested value for the resource.
- D. Kubernetes chooses a random value and uses it as the requested value for the resource.

Answer: C

Explanation:

In Kubernetes, resource management for containers is based on requests and limits. Requests represent the minimum amount of CPU or memory required for scheduling decisions, while limits define the maximum amount a container is

allowed to consume at runtime. Understanding how Kubernetes behaves when only a limit is specified is important for predictable scheduling and resource utilization.

If a container specifies a resource limit but does not explicitly specify a resource request, Kubernetes applies a well-defined default behavior. In this case, Kubernetes automatically sets the request equal to the specified limit. This behavior ensures that the scheduler has a concrete request value to use when deciding where to place the Pod. Without a request value, the scheduler would not be able to make accurate placement decisions, as scheduling is entirely request-based.

This defaulting behavior applies independently to each resource type, such as CPU and memory. For example, if a container sets a memory limit of 512Mi but does not define a memory request, Kubernetes treats the memory request as 512Mi as well. The same applies to CPU limits. As a result, the Pod is scheduled as if it requires the full amount of resources defined by the limit.

Option A is incorrect because specifying only a limit does not cause a container to crash or enter CrashLoopBackOff. CrashLoopBackOff is related to application failures, not resource specification defaults. Option B is incorrect because Kubernetes allows containers to be created without explicit requests, relying on defaulting behavior instead. Option D is incorrect because Kubernetes never assigns random values for resource requests.

This behavior is clearly defined in Kubernetes resource management documentation and is especially relevant when admission controllers like LimitRange are not applying default requests. While valid,

relying solely on limits can reduce cluster efficiency, as Pods may reserve more resources than they actually need. Therefore, best practice is to explicitly define both requests and limits.

Thus, the correct and verified answer is Option C.

Question: 213

In Kubernetes, what is the primary purpose of using annotations?

- A. To control the access permissions for users and service accounts.
- B. To provide a way to attach metadata to objects.
- C. To specify the deployment strategy for applications.
- D. To define the specifications for resource limits and requests.

Answer: B

Explanation:

Annotations in Kubernetes are a flexible mechanism for attaching non-identifying metadata to Kubernetes objects. Their primary purpose is to store additional information that is not used for object selection or grouping, which makes Option B the correct answer.

Unlike labels, which are designed to be used for selection, filtering, and grouping of resources (for example, by Services or Deployments), annotations are intended purely for informational or auxiliary purposes. They allow users, tools, and controllers to store arbitrary key-value data on objects without affecting Kubernetes' core behavior. This makes annotations ideal for storing data such as build information, deployment timestamps, commit hashes, configuration hints, or ownership details.

Annotations are commonly consumed by external tools and controllers rather than by the Kubernetes scheduler or control plane for decision-making. For example, ingress controllers, service meshes, monitoring agents, and CI/CD systems often read annotations to enable or customize specific behaviors. Because annotations are not used for querying or selection, Kubernetes places no strict size or structure requirements on their values beyond general object size limits.

Option A is incorrect because access permissions are managed using Role-Based Access Control (RBAC), which relies on roles, role bindings, and service accounts—not annotations. Option C is incorrect because deployment strategies (such as RollingUpdate or Recreate) are defined in the specification of workload resources like Deployments, not through annotations. Option D is also incorrect because resource limits and requests are specified explicitly in the Pod or container spec under the resources field.

In summary, annotations provide a powerful and extensible way to associate metadata with Kubernetes objects without influencing scheduling, selection, or identity. They support integration, observability, and operational tooling while keeping core Kubernetes behavior predictable and stable. This design intent is clearly documented in Kubernetes metadata concepts, making Option B the correct and verified answer.

Question: 214

What is an important consideration when choosing a base image for a container in a Kubernetes deployment?

- A. It should be minimal and purpose-built for the application to reduce attack surface and improve performance.
- B. It should always be the latest version to ensure access to the newest features.
- C. It should be the largest available image to ensure all dependencies are included.

D. It can be any existing image from the public repository without consideration of its contents.

Answer: A

Explanation:

Choosing an appropriate base image is a critical decision in building containerized applications for Kubernetes, as it directly impacts security, performance, reliability, and operational efficiency. A key best practice is to select a minimal, purpose-built base image, making option A the correct answer.

Minimal base images—such as distroless images or slim variants of common distributions—contain only the essential components required to run the application. By excluding unnecessary packages, shells, and utilities, these images significantly reduce the attack surface. Fewer components mean fewer potential vulnerabilities, which is especially important in Kubernetes environments where containers are often deployed at scale and exposed to dynamic network traffic.

Smaller images also improve performance and efficiency. They reduce image size, leading to faster image pulls, quicker Pod startup times, and lower network and storage overhead. This is particularly beneficial in large clusters or during frequent deployments, scaling events, or rolling updates. Kubernetes' design emphasizes fast, repeatable deployments, and lightweight images align well with these goals.

Option B is incorrect because always using the latest image version can introduce instability or unexpected breaking changes. Kubernetes best practices recommend using explicitly versioned and tested images to ensure predictable behavior and reproducibility. Option C is incorrect because large images increase the attack surface, slow down deployments, and often include unnecessary dependencies that are never used by the application. Option D is incorrect because blindly using public images without inspecting their contents or provenance introduces serious security and compliance risks.

Kubernetes documentation and cloud-native security guidance consistently emphasize the principle of least privilege and minimalism in container images. A well-chosen base image supports secure defaults, faster operations, and easier maintenance, all of which are essential for running reliable workloads in production Kubernetes environments.

Therefore, the correct and verified answer is Option A.

Question: 215

In Kubernetes, what is the primary purpose of creating a Service resource for a Deployment?

- A. To centrally manage and apply runtime configuration values for application components.
- B. To provide a stable endpoint for accessing Pods even when their IP addresses change.
- C. To automatically adjust the number of Pods based on CPU or memory utilization metrics.
- D. To define and attach persistent volumes that store application data across Pod restarts.

Answer: B

Explanation:

In Kubernetes, Pods are inherently ephemeral. They can be created, destroyed, restarted, or rescheduled at any time, and each time this happens, a Pod may receive a new IP address. This dynamic behavior is essential for resilience and scalability, but it also creates a challenge for reliably accessing application workloads. The Service resource addresses this problem by providing a stable network endpoint for a group of Pods, making option B the correct answer.

A Service selects Pods using label selectors—typically the same labels applied by a Deployment—and exposes them through a consistent virtual IP address (ClusterIP) and DNS name. Regardless of how many Pods are running or whether individual Pods are replaced, the Service remains stable and automatically routes traffic to healthy Pods. This abstraction allows clients to communicate with an application without needing to track individual Pod IPs.

Deployments are responsible for managing the lifecycle of Pods, including scaling, rolling updates, and self-healing. However, Deployments do not provide networking or service discovery capabilities. Without a Service, consumers would need to directly reference Pod IPs, which would break as soon as Pods are rescheduled or updated.

Option A is incorrect because centralized configuration management is handled using ConfigMaps and Secrets, not Services. Option C is incorrect because automatic scaling based on CPU or memory is the responsibility of the Horizontal Pod Autoscaler (HPA), not Services. Option D is incorrect because persistent storage is managed using PersistentVolume and PersistentVolumeClaim resources, which are unrelated to Services.

Services can be configured for different access patterns, such as ClusterIP for internal communication, NodePort or LoadBalancer for external access, and headless Services for direct Pod discovery. Despite these variations, their core purpose remains the same: providing a reliable and stable way to access Pods managed by a Deployment.

Therefore, the correct and verified answer is Option B, which aligns with Kubernetes networking fundamentals and official documentation.

Question: 216

A platform engineer is tasked with ensuring that an application can securely access the Kubernetes API without using a developer's personal credentials. What is the correct way to configure this?

- A. Create a ServiceAccount and bind it to the Pod for API access.
- B. Generate a certificate for the application to access the API.
- C. Use a developer's kubeconfig file with restricted permissions.
- D. Set the application to use the default ServiceAccount in the namespace.

Answer: A

Explanation:

In Kubernetes, applications that need to interact with the Kubernetes API should never use a developer's personal credentials. Instead, Kubernetes provides a built-in, secure, and auditable mechanism for workload authentication and authorization using ServiceAccounts. Creating a dedicated ServiceAccount and binding it to the Pod is the correct and recommended approach, making option A the correct answer.

A ServiceAccount represents an identity for processes running inside Pods. When a Pod is configured to use a specific ServiceAccount, Kubernetes automatically injects a short-lived authentication token into the Pod. This token is securely mounted and can be used by the application to authenticate to the Kubernetes API server. Access to API resources is then controlled using RBAC (Role-Based Access Control) by binding roles or cluster roles to the ServiceAccount, ensuring the application has only the permissions it needs—following the principle of least privilege.

Option B is incorrect because manually generating certificates for application access is not the standard or recommended method for in-cluster authentication. Kubernetes manages ServiceAccount tokens automatically and rotates them as needed, providing a simpler and more secure solution. Option C is incorrect because using a developer's kubeconfig file inside an application introduces serious security risks and violates best practices by coupling workloads to personal credentials. Option D is also incorrect because relying on the default ServiceAccount is discouraged; it often has no permissions or, in some cases, broader permissions than intended. Creating a dedicated ServiceAccount provides clearer security boundaries and auditability.

Using ServiceAccounts integrates cleanly with Kubernetes' authentication and authorization model and is explicitly designed for applications and controllers running inside the cluster. This approach ensures secure API access, centralized permission management, and operational consistency across environments.

Therefore, the correct and verified answer is Option A: Create a ServiceAccount and bind it to the Pod for API access.

Question: 217

A request for 500 mebibytes of ephemeral storage must be specified in a YAML file. How should this be written?

- A. 500Mi
- B. 500mi
- C. 500m
- D. 0.5M

Answer: A

Explanation:

In Kubernetes, resource quantities must be expressed using specific, well-defined units. When requesting ephemeral storage, Kubernetes follows the same quantity format rules as other resources such as memory. These rules distinguish between binary units (base-2) and decimal units (base-10), and the correct unit must be used to avoid configuration errors or unintended resource allocation.

The term mebibyte (MiB) is a binary unit equal to 2^{20} bytes (1,048,576 bytes). Kubernetes represents mebibytes using the suffix Mi with a capital "M" and lowercase "i". Therefore, a request for 500 mebibytes of ephemeral storage must be written as 500Mi, making option A the correct answer.

Option B, 500mi, is incorrect because Kubernetes resource units are case-sensitive. The lowercase mi is not a valid unit and will be rejected by the API server. Option C, 500m, is also incorrect because the suffix m represents millicpu when used with CPU resources and has no meaning for storage quantities. Using m for ephemeral storage would result in a validation error. Option D, 0.5M, is incorrect because M represents a decimal megabyte (10^6 bytes), not a mebibyte, and Kubernetes best practices require binary units for memory-based resources like ephemeral storage.

Ephemeral storage requests are typically defined under the container's `resources.requests.ephemeral-storage` field. Correctly specifying the unit ensures that the scheduler can accurately account for node storage capacity and enforce eviction thresholds when necessary.

In summary, Kubernetes requires precise, case-sensitive units for resource specifications. Since the question explicitly asks for 500 mebibytes, the only valid and correct representation is 500Mi, which

aligns exactly with Kubernetes resource quantity conventions.

Question: 218

Which of the following is a primary use case of Istio in a Kubernetes cluster?

- A. To manage and control the versions of container runtimes used on nodes between services.
- B. To provide secure built-in database management features for application workloads.
- C. To provision and manage persistent storage volumes for stateful applications.
- D. To provide service mesh capabilities such as traffic management, observability, and security between services.

Answer: D

Explanation:

Istio is a widely adopted service mesh for Kubernetes that focuses on managing service-to-service communication in distributed, microservices-based architectures. Its primary use case is to provide advanced traffic management, observability, and security capabilities between services, making option D the correct answer.

In a Kubernetes cluster, applications often consist of many independent services that communicate over the network. Managing this communication using application code alone becomes complex and error-prone as systems scale. Istio addresses this challenge by inserting a transparent data plane— typically based on Envoy proxies—alongside application workloads. These proxies intercept all inbound and outbound traffic, enabling consistent policy enforcement without requiring code changes.

Istio's traffic management features include fine-grained routing, retries, timeouts, circuit breaking, fault injection, and canary or blue-green deployments. These capabilities allow operators to control how traffic flows between services, test new versions safely, and improve overall system resilience. For observability, Istio provides detailed telemetry such as metrics, logs, and distributed traces, giving deep insight into service performance and behavior. On the security front, Istio enables mutual TLS (mTLS) for service-to-service communication, strong identity, and access policies to secure traffic within the cluster.

Option A is incorrect because container runtime management is handled at the node and cluster level by Kubernetes and the underlying operating system, not by Istio. Option B is incorrect because Istio does not provide database management functionality. Option C is incorrect because persistent storage provisioning is handled by Kubernetes storage APIs and CSI drivers, not by service meshes.

By abstracting networking concerns away from application code, Istio helps teams operate complex microservices environments more safely and efficiently. Therefore, the correct and verified answer is Option D, which accurately reflects Istio's core purpose and documented use cases in Kubernetes ecosystems.

Question: 219

Kubernetes Secrets are specifically intended to hold confidential data.

a. Which API object should be used to hold non-confidential data?

- A. CNI
- B. CSI
- C. ConfigMaps
- D. RBAC

Answer: C

Explanation:

In Kubernetes, different API objects are designed for different categories of configuration and operational data. Secrets are used to store sensitive information such as passwords, API tokens, and encryption keys. For data that is not confidential, Kubernetes provides the ConfigMap resource, making option C the correct answer.

ConfigMaps are intended to hold non-sensitive configuration data that applications need at runtime. Examples include application configuration files, feature flags, environment-specific settings, URLs, port numbers, and command-line arguments. ConfigMaps allow developers to decouple configuration from application code, which aligns with cloud-native and twelve-factor app principles. This separation makes applications more portable, easier to manage, and simpler to update without

rebuilding container images.

ConfigMaps can be consumed by Pods in several ways: as environment variables, as command-line arguments, or as files mounted into a container's filesystem. Because they are not designed for confidential data, ConfigMaps store values in plaintext and do not provide encryption by default. This is why sensitive data must always be stored in Secrets instead.

Option A, CNI (Container Network Interface), is a networking specification used to configure Pod networking and is

unrelated to data storage. Option B, CSI (Container Storage Interface), is used for integrating external storage systems with Kubernetes and does not store configuration data. Option D, RBAC, defines authorization policies and access controls within the cluster and is not a data storage mechanism.

While both Secrets and ConfigMaps can technically be accessed in similar ways by Pods, Kubernetes clearly distinguishes their intended use cases based on data sensitivity. Using ConfigMaps for non-confidential data improves clarity, security posture, and maintainability of Kubernetes configurations.

Therefore, the correct and verified answer is Option C: ConfigMaps, which are explicitly designed to hold non-confidential configuration data in Kubernetes.

Question: 220

In Kubernetes, what is the primary function of a RoleBinding?

- A. To provide a user or group with permissions across all resources at the cluster level.
- B. To assign the permissions of a Role to a user, group, or service account within a namespace.
- C. To enforce namespace network rules by binding policies to Pods running in the namespace.
- D. To create and define a new Role object that contains a specific set of permissions.

Answer: B

Explanation:

In Kubernetes, authorization is managed using Role-Based Access Control (RBAC), which defines what actions identities can perform on which resources. Within this model, a RoleBinding plays a

crucial role by connecting permissions to identities, making option B the correct answer.

A Role defines a set of permissions—such as the ability to get, list, create, or delete specific resources—but by itself, a Role does not grant those permissions to anyone. A RoleBinding is required to bind that Role to a specific subject, such as a user, group, or service account. This binding is namespace-scoped, meaning it applies only within the namespace where the RoleBinding is created. As a result, RoleBindings enable fine-grained access control within individual namespaces, which is essential for multi-tenant and least-privilege environments.

When a RoleBinding is created, it references a Role (or a ClusterRole) and assigns its permissions to one or more subjects

within that namespace. This allows administrators to reuse existing roles while precisely controlling who can perform certain actions and where. For example, a RoleBinding can grant a service account read-only access to ConfigMaps in a single namespace without affecting access elsewhere in the cluster.

Option A is incorrect because cluster-wide permissions are granted using a ClusterRoleBinding, not a RoleBinding. Option C is incorrect because network rules are enforced using NetworkPolicies, not RBAC objects. Option D is incorrect because Roles are defined independently and only describe permissions; they do not assign them to identities.

In summary, a RoleBinding's primary purpose is to assign the permissions defined in a Role to users, groups, or service accounts within a specific namespace. This separation of permission definition (Role) and permission assignment (RoleBinding) is a fundamental principle of Kubernetes RBAC and is clearly documented in Kubernetes authorization architecture.

Question: 221

A Pod named my-app must be created to run a simple nginx container. Which kubectl command should be used?

- A. `kubectl create nginx --name=my-app`
- B. `kubectl run my-app --image=nginx`
- C. `kubectl create my-app --image=nginx`
- D. `kubectl run nginx --name=my-app`

Answer: B

Explanation:

In Kubernetes, the simplest and most direct way to create a Pod that runs a single container is to use the `kubectl run` command with the appropriate image specification. The command `kubectl run my-app --image=nginx` explicitly instructs Kubernetes to create a Pod named my-app using the nginx container image, which makes option B the correct answer.

The `kubectl run` command is designed to quickly create and run a Pod (or, in some contexts, a higher-level workload resource) from the command line. When no additional flags such as `--restart=Always` are specified, Kubernetes creates a standalone Pod by default. This is ideal for simple use cases like testing, demonstrations, or learning scenarios where only a single container is required.

Option A is incorrect because `kubectl create nginx --name=my-app` is not valid syntax; the `create` subcommand requires a resource type (such as `pod`, `deployment`, or `service`) or a manifest file. Option C is also incorrect because `kubectl create my-app --image=nginx` omits the resource type and therefore is not a valid `kubectl create` command. Option D is incorrect because `kubectl run nginx --name=my-app` attempts to use the deprecated `--name` flag, which is no longer supported in modern versions of `kubectl`.

Using `kubectl run` with explicit naming and image flags is consistent with Kubernetes command-line conventions and is widely documented as the correct approach for creating simple Pods. The resulting Pod can be verified using commands such as `kubectl get pods` and `kubectl describe pod my-app`.

In summary, Option B is the correct and verified answer because it uses valid `kubectl` syntax to create a Pod named `my-app` running the `nginx` container image in a straightforward and predictable way.

Question: 222

What is the primary purpose of a Horizontal Pod Autoscaler (HPA) in Kubernetes?

- A. To automatically scale the number of Pod replicas based on resource utilization.
- B. To track performance metrics and report health status for nodes and Pods.
- C. To coordinate rolling updates of Pods when deploying new application versions.
- D. To allocate and manage persistent volumes required by stateful applications.

Answer: A

Explanation:

The Horizontal Pod Autoscaler (HPA) is a core Kubernetes feature designed to automatically scale the number of Pod replicas in a workload based on observed metrics, making option A the correct answer. Its primary goal is to ensure that applications can handle varying levels of demand while maintaining performance and resource efficiency.

HPA works by continuously monitoring metrics such as CPU utilization, memory usage, or custom and external metrics provided through the Kubernetes metrics APIs. Based on target thresholds defined by the user, the HPA increases or decreases the number of replicas in a scalable resource like a `Deployment`, `ReplicaSet`, or `StatefulSet`. When demand increases, HPA adds more Pods to handle the load. When demand decreases, it scales down Pods to free resources and reduce costs.

Option B is incorrect because tracking performance metrics and reporting health status is handled by components such as the `metrics-server`, monitoring systems, and observability tools—not by the HPA itself. Option C is incorrect because

rolling updates are managed by Deployment strategies, not by the HPA. Option D is incorrect because persistent volume management is handled by Kubernetes storage resources and CSI drivers, not by autoscalers.

HPA operates at the Pod replica level, which is why it is called “horizontal” scaling—scaling out or in by changing the number of Pods, rather than adjusting resource limits of individual Pods (which would be vertical scaling). This makes HPA particularly effective for stateless applications that can scale horizontally to meet demand.

In practice, HPA is commonly used in production Kubernetes environments to maintain application responsiveness under load while optimizing cluster resource usage. It integrates seamlessly with Kubernetes’ declarative model and self-healing mechanisms.

Therefore, the correct and verified answer is Option A, as the Horizontal Pod Autoscaler’s primary function is to automatically scale Pod replicas based on resource utilization and defined metrics.

Question: 223

Which statement best describes the role of kubelet on a Kubernetes worker node?

- A. kubelet manages the container runtime and ensures that all Pods scheduled to the node are running as expected.
- B. kubelet configures networking rules on each node to handle traffic routing for Services in the cluster.
- C. kubelet monitors cluster-wide resource usage and assigns Pods to the most suitable nodes for execution.
- D. kubelet acts as the primary API component that stores and manages cluster state information.

Answer: A

Explanation:

The kubelet is the primary node-level agent in Kubernetes and is responsible for ensuring that workloads assigned to a worker node are executed correctly. Its core function is to manage container execution on the node and ensure that all Pods scheduled to that node are running as expected, which makes option A the correct answer.

Once the Kubernetes scheduler assigns a Pod to a node, the kubelet on that node takes over responsibility for running the Pod. It continuously watches the API server for Pod specifications that target its node and then interacts with the container runtime (such as containerd or CRI-O) through the Container Runtime Interface (CRI). The kubelet starts, stops, and restarts containers to match the desired state defined in the Pod specification.

In addition to lifecycle management, the kubelet performs ongoing health monitoring. It executes liveness, readiness, and startup probes, reports Pod and node status back to the API server, and enforces resource limits defined in the Pod specification. If a container crashes or becomes unhealthy, the kubelet initiates recovery actions such as restarting the container.

Option B is incorrect because configuring Service traffic routing is the responsibility of kube-proxy and the cluster's networking layer, not the kubelet. Option C is incorrect because cluster-wide resource monitoring and Pod placement decisions are handled by the kube-scheduler. Option D is incorrect because cluster state is managed by the API server and stored in etcd, not by the kubelet.

In summary, the kubelet acts as the executor and supervisor of Pods on each worker node. It bridges the Kubernetes control plane and the actual runtime environment, ensuring that containers are running, healthy, and aligned with the declared configuration. Therefore, Option A is the correct and verified answer.

Question: 224

In a cloud native environment, how do containerization and virtualization differ in terms of resource management?

- A. Containerization uses hypervisors to manage resources, while virtualization does not.
- B. Containerization shares the host OS, while virtualization runs a full OS for each instance.
- C. Containerization consumes more memory than virtualization by default.
- D. Containerization allocates resources per container, virtualization does not isolate them.

Answer: B

Explanation:

The fundamental difference between containerization and virtualization in a cloud native environment lies in how they manage and isolate resources, particularly with respect to the operating system. The correct description is that containerization shares the host operating system, while virtualization runs a full operating system for each instance, making option B the correct answer.

In virtualization, each virtual machine (VM) includes its own complete guest operating system running on top of a hypervisor. The hypervisor virtualizes hardware resources—CPU, memory, storage, and networking—and allocates them to each VM. Because every VM runs a full OS, virtualization introduces significant overhead in terms of memory usage, disk space, and startup time. However, it provides strong isolation between workloads, which is useful for running different operating systems or untrusted workloads on the same physical hardware.

In contrast, containerization operates at the operating system level rather than the hardware level. Containers share the host OS kernel and isolate applications using kernel features such as namespaces and control groups (cgroups). This design makes containers much lighter weight than virtual machines. Containers start faster, consume fewer resources, and allow higher workload density on the same infrastructure. Resource limits and isolation are still enforced, but without duplicating the entire operating system for each application instance.

Option A is incorrect because hypervisors are a core component of virtualization, not containerization. Option C is incorrect because containers generally consume less memory than virtual machines due to the absence of a full guest OS. Option D is incorrect because virtualization does isolate resources very strongly, while containers rely on OS-level isolation rather than hardware-level isolation.

In cloud native architectures, containerization is preferred for microservices and scalable workloads because of its efficiency and portability. Virtualization is still valuable for stronger isolation and heterogeneous operating systems.

Therefore, Option B accurately captures the key resource management distinction between the two models.

Question: 225

What is the role of the `ingressClassName` field in a Kubernetes Ingress resource?

- A. It defines the type of protocol (HTTP or HTTPS) that the Ingress Controller should process.
- B. It specifies the backend Service used by the Ingress Controller to route external requests.
- C. It determines how routing rules are prioritized when multiple Ingress objects are applied.
- D. It indicates which Ingress Controller should implement the rules defined in the Ingress resource.

Answer: D

Explanation:

The `ingressClassName` field in a Kubernetes Ingress resource is used to explicitly specify which Ingress Controller is responsible for processing and enforcing the rules defined in that Ingress. This makes option D the correct answer.

In Kubernetes clusters, it is common to have multiple Ingress Controllers running at the same time. For example, a cluster might run an NGINX Ingress Controller, a cloud-provider-specific controller, and an internal-only controller simultaneously. Without a clear mechanism to select which controller should handle a given Ingress resource, multiple controllers could attempt to process the same rules, leading to conflicts or undefined behavior.

The `ingressClassName` field solves this problem by referencing an `IngressClass` object. The `IngressClass` defines the controller implementation (via the `controller` field), and the `Ingress` resource uses `ingressClassName` to declare which class—and therefore which controller—should act on it. This creates a clean and explicit binding between an `Ingress` and its controller.

Option A is incorrect because protocol handling (HTTP vs HTTPS) is defined through TLS configuration and service ports, not by `ingressClassName`. Option B is incorrect because backend Services are

defined in the rules and backend sections of the `Ingress` specification. Option C is incorrect because routing priority is determined by path matching rules and controller-specific logic, not by `ingressClassName`.

Historically, annotations were used to select `Ingress` Controllers, but `ingressClassName` is now the recommended and standardized approach. It improves clarity, portability, and compatibility across different Kubernetes distributions and controllers.

In summary, the primary purpose of `ingressClassName` is to indicate which `Ingress` Controller should implement the routing rules for a given `Ingress` resource, making Option D the correct and verified answer.

Question: 226

In Kubernetes, which command is the most efficient way to check the progress of a `Deployment` rollout and confirm if it has completed successfully?

- A. `kubectl get deployments --show-labels -o wide`
- B. `kubectl describe deployment my-deployment --namespace=default`
- C. `kubectl logs deployment/my-deployment --all-containers=true`
- D. `kubectl rollout status deployment/my-deployment`

Answer: D

Explanation:

When performing rolling updates in Kubernetes, it is important to have a clear and efficient way to track the progress of a `Deployment` rollout and determine whether it has completed successfully. The most direct and purpose-built command for this task is `kubectl rollout status deployment/my-deployment`, making option D the correct answer.

The `kubectl rollout status` command is specifically designed to monitor the state of rollouts for resources such as Deployments, StatefulSets, and DaemonSets. It provides real-time feedback on the rollout process, including whether new Pods have been created, old Pods are being terminated, and if the desired number of updated replicas has become available. The command blocks until the

rollout either completes successfully or fails, which makes it especially useful in automation and CI/CD pipelines.

Option A is incorrect because `kubectl get deployments` only provides a snapshot view of deployment status fields and does not actively track rollout progress. Option B can provide detailed information and events, but it is verbose and not optimized for quickly confirming rollout completion. Option C is incorrect because Deployment objects themselves do not produce logs; logs are generated by Pods and containers, not higher-level workload resources.

The rollout status command also integrates with Kubernetes' revision history, ensuring that it accurately reflects the current state of the Deployment's update strategy. If a rollout is stuck due to failed Pods, readiness probe failures, or resource constraints, the command will indicate that the rollout is not progressing, helping operators quickly identify issues.

In summary, `kubectl rollout status deployment/my-deployment` is the most efficient and reliable way to check rollout progress and confirm success. It is purpose-built for rollout tracking, easy to interpret, and widely used in production Kubernetes workflows, making Option D the correct and verified answer.

Question: 227

When a Kubernetes Secret is created, how is the data stored by default in etcd?

- A. As Base64-encoded strings that provide simple encoding but no actual encryption.
- B. As plain text values that are directly stored without any obfuscation or additional encoding.
- C. As compressed binary objects that are optimized for space but not secured against access.
- D. As encrypted records automatically protected using the Kubernetes control plane master key.

Answer: A

Explanation:

By default, Kubernetes Secrets are stored in etcd as Base64-encoded values, which makes option A the correct answer.

This is a common point of confusion because Base64 encoding is often mistaken for encryption, but in reality, it provides no security—only a reversible text encoding.

When a Secret is defined in a Kubernetes manifest or created via kubectl, its data fields are Base64- encoded before being persisted in etcd. This encoding ensures that binary data (such as certificates or keys) can be safely represented in JSON and YAML formats, which require text-based values. However, anyone with access to etcd or the Secret object via the Kubernetes API can easily decode these values.

Option B is incorrect because Secrets are not stored as raw plaintext; they are encoded using Base64 before storage. Option C is incorrect because Kubernetes does not compress Secret data by default. Option D is incorrect because Secrets are not encrypted at rest by default. Encryption at rest must be explicitly configured using an encryption provider configuration in the Kubernetes API server.

Because of this default behavior, Kubernetes strongly recommends additional security measures when handling Secrets. These include enabling encryption at rest for etcd, restricting access to Secrets using RBAC, using short-lived ServiceAccount tokens, and integrating with external secret management systems such as HashiCorp Vault or cloud provider key management services.

Understanding how Secrets are stored is critical for designing secure Kubernetes clusters. While Secrets provide a convenient abstraction for handling sensitive data, they rely on cluster-level security controls to ensure confidentiality. Without encryption at rest and proper access restrictions, Secret data remains vulnerable to unauthorized access.

Therefore, the correct and verified answer is Option A: Kubernetes stores Secrets as Base64-encoded strings in etcd by default, which offers encoding but not encryption.

Question: 228

In a Kubernetes cluster, which scenario best illustrates the use case for a StatefulSet?

A.

A web application that requires multiple replicas for load balancing.

B.

A service that routes traffic to various microservices in the cluster.

C.

A background job that runs periodically and does not maintain state.

D.

A database that requires persistent storage and stable network identities.

Answer: D

Explanation:

A StatefulSet is a Kubernetes workload API object specifically designed to manage stateful applications. Unlike Deployments or ReplicaSets, which are intended for stateless workloads, StatefulSets provide guarantees about the ordering, uniqueness, and persistence of Pods. These guarantees are critical for applications that rely on stable identities and durable storage, such as databases, message brokers, and distributed systems.

The defining characteristics of a StatefulSet include stable network identities, persistent storage, and ordered deployment and scaling. Each Pod created by a StatefulSet receives a unique and predictable name (for example, database-0, database-1), which remains consistent across Pod restarts. This stable identity is essential for stateful applications that depend on fixed hostnames for leader election, replication, or peer discovery. Additionally, StatefulSets are commonly used with PersistentVolumeClaims, ensuring that each Pod is bound to its own persistent storage that is retained even if the Pod is rescheduled or restarted.

Option A is incorrect because web applications that scale horizontally for load balancing are typically stateless and are best managed by Deployments, which allow Pods to be created and destroyed freely without preserving identity. Option B is incorrect because traffic routing to microservices is handled by Services or Ingress resources, not StatefulSets. Option C is incorrect because periodic background jobs that do not maintain state are better suited for Jobs or CronJobs.

Option D correctly represents the ideal use case for a StatefulSet. Databases require persistent data storage, stable network identities, and predictable startup and shutdown behavior. StatefulSets ensure that Pods are started, stopped, and updated in a controlled order, which helps maintain data consistency and application reliability. According to Kubernetes documentation, whenever an application requires stable identities, ordered deployment, and persistent state, a StatefulSet is the recommended and verified solution, making option D the correct answer.

Question: 229

There is an application running in a logical chain: Gateway API → Service → EndpointSlice → Container.

What Kubernetes API object is missing from this sequence?

- A. Proxy
- B. Docker
- C. Pod
- D. Firewall

Answer: C

Explanation:

In Kubernetes, application traffic flows through a well-defined set of API objects and runtime components before reaching a running container. Understanding this logical chain is essential for grasping how Kubernetes networking works internally.

The given sequence is: Gateway API → Service → EndpointSlice → Container. While this looks close to correct, it is missing a critical Kubernetes abstraction: the Pod. Containers in Kubernetes do not run independently; they always run inside Pods. A Pod is the smallest deployable and schedulable unit in Kubernetes and serves as the execution environment for one or more containers that share networking and storage resources.

The correct logical chain should be:

Gateway API → Service → EndpointSlice → Pod → Container

The Gateway API defines how external or internal traffic enters the cluster. The Service provides a stable virtual IP and DNS name, abstracting a set of backend workloads. EndpointSlices then represent the actual network endpoints backing the Service, typically mapping to the IP addresses of Pods. Finally, traffic is delivered to containers running inside those Pods.

Option A (Proxy) is incorrect because while proxies such as kube-proxy or data plane proxies play a role in traffic forwarding, they are not Kubernetes API objects that represent application workloads in this logical chain. Option B (Docker) is incorrect because Docker is a container runtime, not a Kubernetes API object, and Kubernetes is runtime-agnostic. Option D (Firewall) is incorrect because firewalls are not core Kubernetes workload or networking API objects involved in service-to-container routing.

Option C (Pod) is the correct answer because Pods are the missing link between EndpointSlices and containers.

EndpointSlices point to Pod IPs, and containers cannot exist outside of Pods. Kubernetes documentation clearly states that Pods are the fundamental unit of execution and networking, making them essential in any accurate representation of application traffic flow within a cluster.

Question: 230

During a team meeting, a developer mentions the significance of open collaboration in the cloud native ecosystem. Which statement accurately reflects principles of collaborative development and community stewardship?

A. Open source projects succeed when contributors focus on code quality without the overhead of community engagement.

- B. Maintainers of open source projects act independently to make technical decisions without requiring input from contributors.
- C. Community stewardship emphasizes guiding project growth but does not necessarily include sustainability considerations.
- D. Community events and working groups foster collaboration by bringing people together to share knowledge and build connections.

Answer: D

Explanation:

Open collaboration and community stewardship are foundational principles of the cloud native ecosystem, particularly within projects governed by organizations such as the Cloud Native Computing Foundation (CNCF). These principles emphasize that successful open source projects are not driven solely by code quality, but by healthy, inclusive, and sustainable communities.

Option D accurately reflects these principles. Community events, special interest groups, and working groups play a vital role in fostering collaboration. They provide structured and informal spaces where contributors, maintainers, and users can exchange ideas, share operational experiences, mentor new participants, and collectively guide the direction of projects. This collaborative approach helps ensure that projects evolve in ways that meet real-world needs and benefit from diverse perspectives.

Option A is incorrect because community engagement is not an “overhead” but a critical success factor. Kubernetes and other cloud native projects explicitly recognize that documentation, communication, governance, and contributor onboarding are just as important as writing high- quality code. Without active community participation, projects often struggle with adoption, contributor burnout, and long-term viability.

Option B is incorrect because modern open source governance values transparency and shared decision-making. While maintainers have responsibilities such as reviewing changes and ensuring project stability, they are expected to solicit feedback, encourage discussion, and incorporate contributor input through open processes. This approach builds trust and accountability within the community.

Option C is also incorrect because sustainability is a core aspect of community stewardship. Stewardship includes ensuring that projects can be maintained over time, preventing maintainer burnout, encouraging new contributors, and establishing governance models that support long-term health.

According to cloud native and Kubernetes documentation, strong communities enable innovation, resilience, and scalability—both technically and socially. By bringing people together through events and working groups, community stewardship reinforces collaboration and shared ownership, making option D the correct and fully verified

answer.

E

Question: 231

What does SBOM stand for?

- A. System Bill of Materials
- B. Software Bill Operations Management
- C. Security Baseline for Open Source Management
- D. Software Bill of Materials

Answer: D

Explanation:

SBOM stands for Software Bill of Materials, a critical concept in modern cloud native application delivery and software supply chain security. An SBOM is a formal, structured inventory that lists all components included in a software artifact, such as libraries, frameworks, dependencies, and their versions. This includes both direct and transitive dependencies that are bundled into applications, containers, or container images.

In cloud native environments, applications are often built using numerous open source components and third-party libraries. While this accelerates development, it also increases the risk of hidden vulnerabilities. An SBOM provides transparency into what software is actually running in production, enabling organizations to quickly identify whether they are affected by newly disclosed vulnerabilities or license compliance issues.

Option A is incorrect because SBOM is specific to software, not systems or hardware materials. Option B is incorrect because it describes a management process rather than a standardized inventory of software components. Option C is incorrect because SBOM is not a security baseline or policy framework; instead, it is a factual record of software contents that supports security and compliance efforts.

SBOMs are especially important in containerized and Kubernetes-based workflows. Container images often bundle many dependencies into a single artifact, making it difficult to assess risk without a detailed inventory. By generating and distributing SBOMs alongside container images, teams can integrate vulnerability scanning, compliance checks, and risk

assessment earlier in the delivery pipeline. This practice aligns with the principles of DevSecOps and shift-left security.

Kubernetes and cloud native security guidance emphasize SBOMs as a foundational element of software supply chain security. They support faster incident response, improved trust between software producers and consumers, and stronger governance across the lifecycle of applications. As a result, Software Bill of Materials is the correct and fully verified expansion of SBOM, making option D the accurate answer.

Question: 232

In a Kubernetes cluster, what is the primary role of the Kubernetes scheduler?

- A. To manage the lifecycle of the Pods by restarting them when they fail.
- B. To monitor the health of the nodes and Pods in the cluster.
- C. To handle network traffic between services within the cluster.
- D. To distribute Pods across nodes based on resource availability and constraints.

Answer: D

Explanation:

The Kubernetes scheduler is a core control plane component responsible for deciding where Pods should run within a cluster. Its primary role is to assign newly created Pods that do not yet have a node assigned to an appropriate node based on a variety of factors such as resource availability, scheduling constraints, and policies.

When a Pod is created, it enters a Pending state until the scheduler selects a suitable node. The scheduler evaluates all available nodes and filters out those that do not meet the Pod's requirements. These requirements may include CPU and memory requests, node selectors, node affinity rules, taints and tolerations, topology spread constraints, and other scheduling policies. After filtering, the scheduler scores the remaining nodes to determine the best placement for the Pod and then binds the Pod to the selected node.

Option A is incorrect because restarting failed Pods is handled by other components such as the kubelet and higher-level controllers like Deployments, ReplicaSets, or StatefulSets—not the scheduler. Option B is incorrect because monitoring node and Pod health is primarily the responsibility of the kubelet and the Kubernetes controller manager, which reacts to node failures and ensures desired state. Option C is incorrect because handling network traffic is managed by Services, kube-proxy, and the cluster's networking implementation, not the scheduler.

Option D correctly describes the scheduler's purpose. By distributing Pods across nodes based on resource availability and constraints, the scheduler helps ensure efficient resource utilization, high availability, and workload isolation. This

intelligent placement is essential for maintaining cluster stability and performance, especially in large-scale or multi-tenant environments.

According to Kubernetes documentation, the scheduler's responsibility is strictly focused on Pod placement decisions. Once a Pod is scheduled, the scheduler's job is complete for that Pod, making option D the accurate and fully verified answer.

Question: 233

How does cert-manager integrate with Kubernetes resources to provide TLS certificates for an application?

A.

It manages Certificate resources and Secrets that can be used by Ingress objects for TLS.

B.

It replaces default Kubernetes API certificates with those from external authorities.

C.

It updates kube-proxy configuration to ensure encrypted traffic between Services.

D.

It injects TLS certificates directly into Pods when the workloads are deployed.

Answer: A

Explanation:

cert-manager is a widely adopted Kubernetes add-on that automates the management and lifecycle of TLS certificates in cloud native environments. Its primary function is to issue, renew, and manage certificates by integrating directly with Kubernetes-native resources, rather than modifying core cluster components or injecting certificates manually into workloads.

Option A correctly describes how cert-manager operates. cert-manager introduces Custom Resource Definitions (CRDs) such as Certificate, Issuer, and ClusterIssuer. These resources define how certificates should be requested and from which certificate authority they should be obtained, such as Let's Encrypt or a private PKI. Once a certificate is successfully issued, cert-manager stores it in a Kubernetes Secret. These Secrets can then be referenced by Ingress resources, Gateway API resources, or directly by applications to enable TLS.

Option B is incorrect because cert-manager does not replace or interfere with Kubernetes API server certificates. The Kubernetes control plane manages its own internal certificates independently, and cert-manager is focused on application-level TLS, not control plane security.

Option C is incorrect because cert-manager does not interact with kube-proxy or manage service-to-service encryption. Traffic encryption between Services is typically handled by service meshes or application-level TLS configurations, not cert-manager.

Option D is incorrect because cert-manager does not inject certificates directly into Pods at deployment time. Instead, Pods consume certificates indirectly by mounting the Secrets created and maintained by cert-manager. This design aligns with Kubernetes best practices by keeping certificate management decoupled from application deployment logic.

According to Kubernetes and cert-manager documentation, cert-manager's strength lies in its native integration with Kubernetes APIs and declarative workflows. By managing Certificate resources and automatically maintaining Secrets for use by Ingress or Gateway resources, cert-manager simplifies TLS management, reduces operational overhead, and improves security across cloud native application delivery pipelines. This makes option A the accurate and fully verified answer.

Question: 234

A Pod has been created, but when checked with `kubectl get pods`, the READY column shows 0/1.

What Kubernetes feature causes this behavior?

- A. Node Selector
- B. Readiness Probes
- C. DNS Policy
- D. Security Contexts

Answer: B

Explanation:

The READY column in the output of `kubectl get pods` indicates how many containers in a Pod are currently considered ready to serve traffic, compared to the total number of containers defined in that Pod. A value of 0/1 means that the Pod has one container, but that container is not yet marked as ready. The Kubernetes feature responsible for determining this readiness state is the readiness probe.

Readiness probes are used by Kubernetes to decide when a container is ready to accept traffic. These probes can be configured to perform HTTP requests, execute commands, or check TCP sockets inside the container. If a readiness probe is defined and it fails, Kubernetes marks the container as not ready, even if the container is running successfully. As a

result, the READY column will show 0/1, and the Pod will be excluded from Service load balancing until the probe succeeds.

Option A (Node Selector) is incorrect because node selectors influence where a Pod is scheduled, not whether its containers are considered ready after startup. Option C (DNS Policy) affects how DNS resolution works inside a Pod and has no direct impact on readiness reporting. Option D (Security Contexts) define security-related settings such as user IDs, capabilities, or privilege levels, but they do not control the READY status shown by kubectl.

Readiness probes are particularly important for applications that take time to initialize, load configuration, or warm up caches. By using readiness probes, Kubernetes ensures that traffic is only sent to containers that are fully prepared to handle requests. This improves reliability and prevents failed or premature connections.

According to Kubernetes documentation, a container without a readiness probe is considered ready by default once it is running. However, when a readiness probe is defined, its result directly controls the READY state. Therefore, the presence and behavior of readiness probes is the verified and correct

reason why a Pod may show 0/1 in the READY column, making option B the correct answer.

Question: 235

What does the livenessProbe in Kubernetes help detect?

- A. When a container is ready to serve traffic.
- B. When a container has started successfully.
- C. When a container exceeds resource limits.
- D. When a container is unresponsive.

Answer: D

Explanation:

The liveness probe in Kubernetes is designed to detect whether a container is still running correctly or has entered a failed or unresponsive state. Its primary purpose is to determine whether a container should be restarted. When a liveness probe fails repeatedly, Kubernetes assumes the container is unhealthy and automatically restarts it to restore normal operation.

Option D correctly describes this behavior. Liveness probes are used to identify situations where an application is running but no longer functioning as expected—for example, a deadlock, infinite loop, or hung process that cannot recover on its own. In such cases, restarting the container is often the most effective remediation, and Kubernetes handles this automatically through the liveness probe mechanism.

Option A is incorrect because readiness probes—not liveness probes—determine whether a container is ready to receive traffic. A container can be alive but not ready, such as during startup or temporary maintenance. Option B is incorrect because startup success is handled by startup probes, which are specifically designed to manage slow-starting applications and delay liveness and readiness checks until initialization is complete. Option C is incorrect because exceeding resource limits is managed by the container runtime and kubelet (for example, OOMKills), not by probes.

Liveness probes can be implemented using HTTP requests, TCP socket checks, or command execution inside the container. If the probe fails beyond a configured threshold, Kubernetes restarts the container according to the Pod's restart policy. This self-healing behavior is a core feature of Kubernetes and contributes significantly to application reliability.

Kubernetes documentation emphasizes using liveness probes carefully, as misconfiguration can cause unnecessary restarts. However, when used correctly, they provide a powerful way to automatically recover from application-level failures that Kubernetes cannot otherwise detect.

In summary, the liveness probe's role is to detect when a container is unresponsive and needs to be restarted, making option D the correct and fully verified answer.

Question: 236

When modifying an existing Helm release to apply new configuration values, which approach is the best practice?

- A. Use helm upgrade with the --set flag to apply new values while preserving the release history.
- B. Use kubectl edit to modify the live release configuration and apply the updated resource values.
- C. Delete the release and reinstall it with the desired configuration to force an updated deployment.
- D. Edit the Helm chart source files directly and reapply them to push the updated configuration values.

Answer: A

Explanation:

Helm is a package manager for Kubernetes that provides a declarative and versioned approach to application deployment and lifecycle management. When updating configuration values for an existing Helm release, the recommended and best-practice approach is to use helm upgrade, optionally with the --set flag or a values file, to apply the new configuration while preserving the release's history.

Option A is correct because helm upgrade updates an existing release in a controlled and auditable manner. Helm stores each revision of a release, allowing teams to inspect past configurations and roll back to a previous known-good state if needed. Using --set enables quick overrides of individual values, while using -f values.yaml supports more complex or repeatable configurations. This approach aligns with GitOps and infrastructure-as-code principles, ensuring consistency and traceability.

Option B is incorrect because modifying Helm-managed resources directly with kubectl edit breaks Helm's state tracking. Helm maintains a record of the desired state for each release, and manual edits can cause configuration drift, making future upgrades unpredictable or unsafe. Kubernetes documentation and Helm guidance strongly discourage modifying Helm-managed resources outside of Helm itself.

Option C is incorrect because deleting and reinstalling a release discards the release history and may cause unnecessary downtime or data loss, especially for stateful applications. Helm's upgrade mechanism is specifically designed to avoid this disruption while still applying configuration changes

safely.

Option D is also incorrect because editing chart source files directly and reapplying them bypasses Helm's release management model. While chart changes are appropriate during development, applying them directly to a running release without helm upgrade undermines versioning, rollback, and repeatability.

According to Helm documentation, helm upgrade is the standard and supported method for modifying deployed applications. It ensures controlled updates, preserves operational history, and enables safe rollbacks, making option A the correct and fully verified best practice.

Question: 237

Which Kubernetes Service type exposes a service only within the cluster?

A.ClusterIP

B.NodePort

C.LoadBalancer

D.ExternalName

Answer: A

Explanation:

In Kubernetes, a Service provides a stable network endpoint for a set of Pods and abstracts away their dynamic nature. Kubernetes offers several Service types, each designed for different exposure requirements. Among these, ClusterIP is the Service type that exposes an application only within the cluster, making it the correct answer.

When a Service is created with the ClusterIP type, Kubernetes assigns it a virtual IP address that is reachable exclusively from within the cluster's network. This IP is used by other Pods and internal components to communicate with the Service through cluster DNS or environment variables. External traffic from outside the cluster cannot directly access a ClusterIP Service, which makes it ideal for internal APIs, backend services, and microservices that should not be publicly exposed.

Option B (NodePort) is incorrect because NodePort exposes the Service on a static port on each node's IP address,

allowing access from outside the cluster. Option C (LoadBalancer) is incorrect because it provisions an external load balancer—typically through a cloud provider—to expose the Service publicly. Option D (ExternalName) is incorrect because it does not create a proxy or internal endpoint at all; instead, it maps the Service name to an external DNS name outside the cluster.

ClusterIP is also the default Service type in Kubernetes. If no type is explicitly specified in a Service

manifest, Kubernetes automatically assigns it as ClusterIP. This default behavior reflects the principle of least exposure, encouraging internal-only access unless external access is explicitly required.

From a cloud native architecture perspective, ClusterIP Services are fundamental to building secure, scalable microservices systems. They enable internal service-to-service communication while reducing the attack surface by preventing unintended external access.

According to Kubernetes documentation, ClusterIP Services are intended for internal communication within the cluster and are not reachable from outside the cluster network. Therefore, ClusterIP is the correct and fully verified answer, making option A the right choice.

Question: 238

Which option best represents the Pod Security Standards ordered from most permissive to most restrictive?

A.

Privileged, Baseline, Restricted

B.

Baseline, Privileged, Restricted

C.

Baseline, Restricted, Privileged

D.

Privileged, Restricted, Baseline

Answer: A

Explanation:

Pod Security Standards define a set of security profiles for Pods in Kubernetes, establishing clear expectations for how securely workloads should be configured. These standards were introduced to replace the deprecated PodSecurityPolicies (PSP) and are enforced through the Pod Security Admission controller. The standards are intentionally ordered from least restrictive to most restrictive to allow clusters to adopt security controls progressively.

The correct order from most permissive to most restrictive is: Privileged → Baseline → Restricted, which makes option A the correct answer.

The Privileged profile is the least restrictive. It allows Pods to run with elevated permissions, including privileged containers, host networking, host PID/IPC namespaces, and unrestricted access to host resources. This level is intended for trusted system components, infrastructure workloads, or cases where full access to the host is required. It offers maximum flexibility but minimal security enforcement.

The Baseline profile introduces a moderate level of security. It prevents common privilege escalation vectors, such as running privileged containers or using host namespaces, while still allowing typical application workloads to function without significant modification. Baseline is designed to be broadly compatible with most applications and serves as a reasonable default security posture for many clusters.

The Restricted profile is the most secure and restrictive. It enforces strong security best practices, such as requiring containers to run as non-root users, dropping unnecessary Linux capabilities, enforcing read-only root filesystems where possible, and preventing privilege escalation. Restricted is ideal for highly sensitive workloads or environments with strict security requirements, though it may require application changes to comply.

Options B, C, and D are incorrect because they misrepresent the intended progression of security strictness defined in Kubernetes documentation.

According to Kubernetes documentation, the Pod Security Standards are explicitly ordered to support gradual adoption: start permissive where necessary and move toward stronger security over time. Therefore, Privileged, Baseline, Restricted is the accurate and fully verified ordering, making option A the correct answer.

Question: 239

Which option represents best practices when building container images?

- A. Use multi-stage builds, use the latest tag for image version, and only install necessary packages.
- B. Use multi-stage builds, pin the base image version to a specific digest, and install extra packages just in case.
- C. Use multi-stage builds, pin the base image version to a specific digest, and only install necessary packages.
- D. Avoid multi-stage builds, use the latest tag for image version, and install extra packages just in case.

Answer: C

Explanation:

Building secure, efficient, and reproducible container images is a core principle of cloud native application delivery. Kubernetes documentation and container security best practices emphasize minimizing image size, reducing attack surface, and ensuring deterministic builds. Option C fully aligns with these principles, making it the correct answer.

Multi-stage builds allow developers to separate the build environment from the runtime environment. Dependencies such as compilers, build tools, and temporary artifacts are used only in intermediate stages and excluded from the final image. This significantly reduces image size and limits the presence of unnecessary tools that could be exploited at runtime.

Pinning the base image to a specific digest ensures immutability and reproducibility. Tags such as latest can change over time, potentially introducing breaking changes or vulnerabilities without notice. By using a digest, teams guarantee that the same base image is used every time the image is built, which is essential for predictable behavior, security auditing, and reliable rollbacks.

Installing only necessary packages further reduces the attack surface. Every additional package increases the risk of vulnerabilities and expands the maintenance burden. Minimal images are faster to pull, quicker to start, and easier to scan for vulnerabilities. Kubernetes security guidance consistently recommends keeping container images as small and purpose-built as possible.

Option A is incorrect because using the latest tag undermines build determinism and traceability. Option B is incorrect because installing extra packages "just in case" contradicts the principle of minimalism and increases security risk. Option D is incorrect because avoiding multi-stage builds and installing unnecessary packages leads to larger, less secure images and is explicitly discouraged in cloud native best practices.

According to Kubernetes and CNCF security guidance, combining multi-stage builds, immutable image references, and minimal dependencies results in more secure, reliable, and maintainable container images. Therefore, option C represents the best and fully verified approach when building container images.

Question: 240

A Kubernetes Pod is returning a CrashLoopBackOff status. What is the most likely reason for this behavior?

- A. There are insufficient resources allocated for the Pod.
- B. The application inside the container crashed after starting.
- C. The container's image is missing or cannot be pulled.
- D. The Pod is unable to communicate with the Kubernetes API server.

s **Answer: B**

Explanation:

A CrashLoopBackOff status in Kubernetes indicates that a container within a Pod is repeatedly starting, crashing, and

being restarted by Kubernetes. This behavior occurs when the container process exits shortly after starting and Kubernetes applies an increasing back-off delay between restart attempts to prevent excessive restarts.

Option B is the correct answer because CrashLoopBackOff most commonly occurs when the application inside the container crashes after it has started. Typical causes include application runtime errors, misconfigured environment variables, missing configuration files, invalid command or entrypoint definitions, failed dependencies, or unhandled exceptions during application startup. Kubernetes itself is functioning as expected by restarting the container according to the Pod's restart policy.

Option A is incorrect because insufficient resources usually lead to different symptoms. For example, if a container exceeds its memory limit, it may be terminated with an OOMKilled status rather than repeatedly crashing immediately.

While resource constraints can indirectly cause crashes, they are not the defining reason for a CrashLoopBackOff state.

Option C is incorrect because an image that cannot be pulled results in statuses such as ImagePullBackOff or ErrImagePull, not CrashLoopBackOff. In those cases, the container never successfully starts.

Option D is incorrect because Pods do not need to communicate directly with the Kubernetes API server for normal application execution. Issues with API server communication affect control plane components or scheduling, not container restart behavior.

From a troubleshooting perspective, Kubernetes documentation recommends inspecting container logs using `kubectl logs` and reviewing Pod events with `kubectl describe pod` to identify the root cause of the crash. Fixing the underlying application error typically resolves the CrashLoopBackOff condition.

In summary, CrashLoopBackOff is a protective mechanism that signals a repeatedly failing container process. The most likely and verified cause is that the application inside the container is crashing after startup, making option B the correct answer.