LinuxFoundation.CGOA.v2026-04-11.q20

Exam Code:	CGOA
Exam Name:	Certified GitOps Associate
Certification Provider:	Linux Foundation
Free Question Number:	20
Version:	v2026-04-11
# of views:	107
# of Questions views:	200
https://www.freeqas.com/qa/Linux-Foundation/CGOA/LinuxFoundation.CGOA.v2026-04-11.q20.html	

NEW QUESTION: 1
Which two GitOps controllers are the most popular?
A. Keptn and Puppet
B. Helm and Kustomize
C. ArgoCD and Flux
D. Jenkins and Tekton
Answer: C (LEAVE A REPLY)
The two most widely adopted GitOps controllers in the CNCF ecosystem areArgoCDandFlux. Both implement the GitOps principles of continuous reconciliation from Git as the source of truth.
"ArgoCD and Flux are the two primary CNCF GitOps controllers. They implement reconciliation loops to ensure the desired state in Git matches the cluster's actual state."
Thus, the correct answer isC.
References:GitOps Tooling (CNCF GitOps Working Group).

NEW QUESTION: 2
What is the main difference between Terraform/OpenTofu and Ansible?
A. Terraform/OpenTofu uses a configuration language called CUE, while Ansible uses HCL.
B. Terraform/OpenTofu stores the state of each resource, while Ansible works in a fire-and-forget mode.
C. Terraform/OpenTofu is imperative in nature, while Ansible is declarative.
D. Ansible is written in Golang, while Terraform/OpenTofu is written in Python.
Answer: B (LEAVE A REPLY)
Terraform (or OpenTofu)uses a declarative model and maintains astate fileto track the current status of resources, enabling it to plan and reconcile changes.Ansible, by contrast,

is more procedural and executes tasks in afire-and-forgetmanner, without tracking persistent resource state.
"Terraform maintains state for each managed resource, enabling planned, consistent changes. Ansible executes tasks without tracking resource state, working in a fire-and-forget model." Thus, the correct answer isB.
References:GitOps Tooling (CNCF GitOps Working Group).

NEW QUESTION: 3
You want to create a dashboard to monitor the performance of your application. Which of the following is a key principle of GitOps regarding dashboards?
A. The operations team should manually update dashboards.
B. Dashboards should be created using a proprietary tool.
C. Dashboards should only be accessible to the development team.
D. Dashboards declarations should be in the Desired State store.
Answer: (SHOW ANSWER)
In GitOps,everything that defines the system, including dashboards, must be stored declaratively in Git(the Desired State store). This ensures dashboards are versioned, reproducible, and consistent across environments.
"GitOps requires that all system components, including monitoring and observability configurations such as dashboards, are declared in Git. This ensures they are versioned, immutable, and reproducible." Thus,Dis correct.
References:GitOps Principles (CNCF GitOps Working Group).

NEW QUESTION: 4
What is one of the key benefits of a pull-based reconciliation approach to configuration management?
A. Simplified troubleshooting and debugging processes.
B. Agents can access the Desired State at any time, not only when an event is triggered.
C. The CI has access credentials to the running system.
D. Immediate response time to configuration changes.
Answer: B (LEAVE A REPLY)
In GitOps, thepull-based reconciliation approachmeans that agents continuously monitor the Desired State in Git. Unlike push-based systems, which only act when triggered, pull-based systems can reconcile at any time, providing resilience, self-healing, and security (since no external system needs direct access to the cluster).
"In a pull-based model, reconciliation agents continuously fetch and compare the desired state, enabling self- healing and ensuring the desired configuration is accessible at all times." Thus, the correct answer isB.
References:GitOps Principles (CNCF GitOps Working Group), Pull vs. Push reconciliation models.

NEW QUESTION: 5
How can you achieve the declarative GitOps principle in managing infrastructure and applications?
A. By using imperative scripting languages to automate infrastructure changes.
B. By manually making ad-hoc configuration changes directly in the production environment.
C. By periodically creating manual backups of your infrastructure configurations.
D. By defining and maintaining infrastructure and application configurations declaratively in a version- controlled system.
Answer: (SHOW ANSWER)
The first GitOps principle isDeclarative Descriptions. This means the desired system configuration (for infrastructure, services, and applications) is expressed declaratively and stored in version control. Git becomes the single source of truth.
"The desired system state must be expressed declaratively. This provides a clear, machine-readable blueprint for the system, and ensures that what is in Git is what should be running in the environment." Therefore,infrastructure and application configurations must be defined declaratively and stored in Git, not managed imperatively or manually.
References:GitOps Principles (CNCF GitOps Working Group),Principle 1: The system is described declaratively.

NEW QUESTION: 6
Which statement describes Blue-Green deployments?
A. Blue-Green deployments involve deploying the new version of an application alongside the old version and switching traffic to the latest version once it is ready.
B. Blue-Green deployments involve deploying the new version of an application to a subset of users and gradually expanding the deployment based on feedback.
C. Blue-Green deployments involve deploying different versions of an application in other regions and routing traffic based on geographic location.
D. Blue-Green deployments involve deploying only one version at a time.
Answer: A (LEAVE A REPLY)
Blue-Green deployments are aprogressive delivery patternwhere two environments exist: Blue (current version) and Green (new version). The new version is deployed in parallel, and once validated, traffic is switched over from Blue to Green.
"Blue-Green deployments provide zero-downtime releases by running two production environments: one active and one idle. A new version is deployed to the idle environment, tested, and when ready, traffic is switched to it." Thus, the correct description isA.
References:GitOps Patterns (CNCF GitOps Working Group), Progressive Delivery patterns.

NEW QUESTION: 7

Which deployment and release pattern involves gradually shifting traffic from an old version of an application to a new one?

A. Red/Black Deployment

B. Canary Deployment

C. Blue-Green Deployment

D. A/B Deployment

Answer: B (LEAVE A REPLY)

ACanary Deploymentgradually introduces a new application version to a small subset of users before expanding to the full user base. This pattern allows testing and validation in production while reducing risk.

"Canary deployments progressively roll out changes to a small group of users, monitoring for issues before routing all traffic to the new version. This gradual shift minimizes risk and ensures safer releases." Thus, the correct answer isB.

References:GitOps Patterns (CNCF GitOps Working Group), Progressive Delivery.

NEW QUESTION: 8

What doesPulled Automaticallyrefer to?

A. A GET request to a relational database.

B. It always refers to Git pull.

C. Webhooks informing the system about new commits.

D. Accessing the Desired State from the State Store.

Answer: D (LEAVE A REPLY)

ThePulled AutomaticallyGitOps principle refers to the way software agents continuously access theDesired Statestored in the State Store (e.g., Git). Agents automatically pull the state from the repository and reconcile the system accordingly.

"Software agents automatically pull the desired state declarations from the source of truth (State Store) and continuously reconcile the system to match." Thus, the correct answer isD.

References:GitOps Principles (CNCF GitOps Working Group).

NEW QUESTION: 9

You are working on a GitOps project and need to understand the similarities and differences between pull- based messaging systems and event-driven systems. What is a key difference between these two types of systems?

A. Pull-based systems require a constant network connection to receive updates.

B. Event-driven systems are less flexible and scalable compared to pull-based systems.

C. Pull-based systems are more efficient in handling real-time events.

D. When only events trigger reconciliation, the system is more vulnerable to drift caused by other things.

Answer: D (LEAVE A REPLY)

In GitOps, thepull-based modelcontinuously reconciles the actual state with the desired state. This makes it resilient to drift, since reconciliation runs regularly. In contrast,event-driven systemsonly reconcile when an event occurs (e.g., a webhook), which makes them more prone to drift if changes happen outside those events.
"A pull-based reconciliation loop ensures continuous alignment with the desired state. Event-driven reconciliation, triggered only on events, risks system drift if changes occur outside those triggers." Thus, the correct answer isD.
References:GitOps Related Practices (CNCF GitOps Working Group), Reconciliation Models.

NEW QUESTION: 10
In the context of GitOps, why would you do a rollback?
A. To undo a deployment that introduced a critical bug or caused a system failure.
B. To improve performance and optimize resource utilization.
C. To create a backup of the current configuration.
D. To test a new feature in a controlled environment.
Answer: (SHOW ANSWER)
In GitOps, rollback is the process of reverting to a previous known-good configuration stored in Git. This is typically done when a deployment introduces abug, error, or failurethat impacts system stability.
"Rollback in GitOps is used to revert to a previous commit representing a stable configuration when the current deployment causes errors or failures." Thus, the correct answer isA.
References:GitOps Principles (CNCF GitOps Working Group), Rollback and Recovery.

NEW QUESTION: 11
In GitOps, how is the Desired State stored?
A. In a way that enforces mutability and versioning.
B. In a way that permits direct modifications to live systems.
C. In a way that retains only the latest version.
D. In a way that enforces immutability and versioning.
Answer: D (LEAVE A REPLY)
The GitOps principle ofVersioned and Immutablerequires Desired State to be stored in a way that enforcesimmutability and versioning. This ensures every change is recorded, auditable, and reversible.
"Desired state must be kept in an immutable, version-controlled system. This guarantees a full history of changes and enables safe rollbacks." Thus, the correct answer isD.
References:GitOps Principles (CNCF GitOps Working Group).

NEW QUESTION: 12

You are working on a GitOps project and want to trigger a reconcile process before the next scheduled reconciliation. What is the recommended way to do this?
A. Adjust the reconcile process interval time.
B. Schedule a cron job to run the reconcile process periodically, using RBAC to authenticate.
C. Manually execute a script to initiate the reconcile process on the cluster using GitOps tool CLI commands.
D. Use a webhook to trigger the reconcile process based on events or changes in the Git repository.
Answer: D (LEAVE A REPLY)
Although reconciliation is continuous in GitOps, tools often allow reconciliation to be triggered earlier than the normal polling interval. The recommended approach is to usewebhooks from the Git repository, which notify the GitOps controller of changes and trigger an immediate reconcile.
"While reconciliation loops continuously compare desired and actual state, reconciliation can be triggered sooner by webhooks from version control events, ensuring timely application of changes." Thus, the correct answer isA.
References:GitOps Principles (CNCF GitOps Working Group), Reconciliation and Webhooks.

NEW QUESTION: 13
In the context of GitOps, what source of truth guides the continuous deployment process?
A. Desired State
B. Dynamic State
C. Fleeting State
D. Current State
Answer: A (LEAVE A REPLY)
TheDesired State, stored in Git, is the ultimate source of truth in GitOps. It defines how the system should look and behave. Continuous deployment processes reconcile the actual cluster state against this Desired State.
"In GitOps, the desired state kept in Git is the single source of truth. The reconciler ensures the actual state matches the desired state, guiding the continuous deployment process." Thus, the correct answer isA.
References:GitOps Terminology (CNCF GitOps Working Group).

NEW QUESTION: 14
You are deploying a new version of your application using the Blue-Green deployment pattern. What is a characteristic of the Blue-Green deployment pattern?
A. The new version of the application is deployed first, followed by the old version.
B. The old version of the application is deployed first, followed by the new version.
C. Both the new and old versions of the application are deployed simultaneously.

D. The Blue-Green deployment pattern only deploys single versions of the application.
Answer: (SHOW ANSWER)
In aBlue-Green deployment, two environments (Blue and Green) exist at the same time. The current version runs in one environment (Blue), and the new version is deployed to the other environment (Green). Traffic is switched to Green once the new version is validated. "Blue-Green deployments maintain two production environments. The new version is deployed alongside the old version, and once validated, traffic is switched from Blue to Green." Thus, the correct answer isC.
References:GitOps Patterns (CNCF GitOps Working Group), Progressive Delivery.

NEW QUESTION: 15
In GitOps practices, when does CD take part?
A. CD takes part simultaneously with CI, both components of GitOps practices.
B. CD takes part after CI to automate the deployment of applications based on changes in the Git repository.
C. CD takes part before CI stage in order to ensure the successful deployment of applications.
D. CI plays a significant role in GitOps practices.
Answer: B (LEAVE A REPLY)
In GitOps, Continuous Deployment (CD) follows after Continuous Integration (CI). CI is responsible for building and testing application code, while CD automates the delivery and deployment of these changes into runtime environments. The Git repository serves as the single source of truth, and when CI merges new changes into the main branch, CD reconciles the state of the environment to match what is declared in Git.
"GitOps builds on the principles of DevOps by using Git as the source of truth for declarative infrastructure and applications. CI pipelines handle the integration and testing of code, and CD pipelines or agents automatically reconcile the desired state in Git with the actual state in the cluster." This shows that CD is triggered after CI to handle deployment automation, ensuring systems remain in sync with what is declared in version control.
References: GitOps Principles (CNCF GitOps Working Group), GitOps Working Group Terminology & Principles documents.

NEW QUESTION: 16
In GitOps, what is the process of ensuring the actual state of a system matches its Desired State called?
A. Reconciliation
B. Webhooks
C. Monitoring
D. Manual Intervention
Answer: (SHOW ANSWER)

The process of keeping theactual statein sync with thedesired stateis calledReconciliation. GitOps controllers (e.g., ArgoCD, Flux) continuously reconcile system resources to match what is declared in Git.
"Reconciliation is the process by which agents compare the actual system state to the desired state and automatically make changes to converge them." Thus, the correct answer isA: Reconciliation.
References:GitOps Terminology (CNCF GitOps Working Group).

NEW QUESTION: 17
You are implementing GitOps in your organization and have configured the Desired State of your applications in a Git repository. However, during the deployment process, you encounter an error in the configuration. What is the recommended action in this scenario?
A. Continue to monitor the issue and proceed with the deployment, as it may not significantly impact the application.
B. Raise a ticket with the development team to fix the error in the configuration file.
C. Roll back the deployment to the previous working version while investigating the error in the configuration file.
D. Make a call to the Kubernetes API with the correction.
Answer: C (LEAVE A REPLY)
GitOps emphasizes immutability and auditability. If an error occurs in the configuration stored in Git, the system should be rolled back to the last known good state while the error is fixed. This preserves system reliability and aligns with the GitOps principle of rollback through version control.
"With Git as the source of truth, if an error is introduced, the system can be rolled back by reverting to a previous commit. This ensures stability while the faulty configuration is corrected." Thus, the recommended action isC: Roll back to the previous working version.
References:GitOps Principles (CNCF GitOps Working Group).

NEW QUESTION: 18
You are working on a GitOps deployment and want to manage the configuration of your Kubernetes resources across multiple environments. How does Kustomize help?
A. Kustomize is a tool for deploying infrastructure resources using Terraform/OpenTofu.

B. Kustomize allows you to package and distribute your application as a Helm chart.
C. Kustomize helps you create and manage Kubernetes resource manifests by providing a graphical user interface (GUI).
D. Kustomize helps you create and manage Kubernetes resource manifests by providing a way to customize them through patching.
Answer: (SHOW ANSWER)
Kustomizeis a Kubernetes-native configuration management tool that allows manifest customization without modifying the original YAML files. It uses overlays and patches to adapt configurations for different environments.
"Kustomize provides a declarative way to customize Kubernetes manifests by applying patches and overlays.
This allows managing multiple environments without duplicating manifest files." Thus, the correct answer isD.
References:GitOps Tooling (CNCF GitOps Working Group), Kustomize.

NEW QUESTION: 19
In the context of GitOps, what is one example of how DevSecOps principles manifested, enhancing the traditional DevOps lifecycle?
A. GitOps enhances the DevSecOps experience by detecting security policy drift.
B. DevSecOps in GitOps focuses primarily on post-deployment security audits.
C. GitOps uses DevSecOps to enforce manual security checks at each deployment stage.
D. In GitOps, DevSecOps leads to the segregation of security tasks, assigning them exclusively to security teams.
Answer: (SHOW ANSWER)
In GitOps,DevSecOpsintegrates security into the GitOps workflow by treating security policies as code and storing them in Git. This enables automatic detection ofsecurity policy driftand ensures that any misconfiguration or violation is reconciled, just like application and infrastructure code.
"GitOps applies DevSecOps by managing security policies as code. This enables detection of drift in security configurations, ensuring environments remain compliant and secure." Thus, the correct answer isA.
References:GitOps Related Practices (CNCF GitOps Working Group), DevSecOps integration.

NEW QUESTION: 20
In the context of GitOps, what happens to a GitOps-managed Kubernetes cluster if there is drift divergence?
A. The GitOps-managed Kubernetes cluster ignores the drift divergence and continues to operate as it is.
B. The GitOps-managed Kubernetes cluster automatically reconciles the drift divergence to return the cluster to the Desired State.

C. The GitOps-managed Kubernetes cluster notifies the administrator about the drift divergence and waits for manual intervention.
D. The GitOps-managed Kubernetes cluster rolls back to the previous known state before the drift divergence occurred.
Answer: (SHOW ANSWER)
A GitOps-managed Kubernetes cluster uses reconciliation loops to continuously compare the actual state of the system with the desired state declared in Git. When drift (divergence between declared configuration and live cluster state) is detected, the GitOps operator automatically reconciles the difference to bring the system back into alignment.
"In GitOps, a reconciliation loop ensures that the desired state as declared in Git is continuously compared with the observed state of the system. If drift is detected, the system automatically takes corrective action to reconcile the difference and restore the declared configuration." This ensures consistency, reliability, and self-healing. Manual intervention is not required for drift correction, as the automated reconciliation is a core principle of GitOps.
References: GitOps Principles (CNCF GitOps Working Group), GitOps Principles Document -Principle 4:
Software agents automatically pull the desired state declarations from the source and continuously observe actual system state, reconciling differences.