PROFESSIONAL-CLOUD-DEVELOPER Q&As

Professional Cloud Developer

Pass Google PROFESSIONAL-CLOUD-DEVELOPER Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.leads4pass.com/professional-cloud-developer.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google Official Exam Center

QUESTION 1

Your team is developing an ecommerce platform for your company. Users will log in to the website and add items to their shopping cart. Users will be automatically logged out after 30 minutes of inactivity. When users log back in, their shopping cart should be saved. How should you store users\\' session and shopping cart information while following Google- recommended best practices?

A. Store the session information in Pub/Sub, and store the shopping cart information in Cloud SQL.

B. Store the shopping cart information in a file on Cloud Storage where the filename is the SESSION ID.

C. Store the session and shopping cart information in a MySQL database running on multiple Compute Engine instances.

D. Store the session information in Memorystore for Redis or Memorystore for Memcached, and store the shopping cart information in Firestore.

Correct Answer: D

QUESTION 2

You are writing a single-page web application with a user-interface that communicates with a third-party API

for content using XMLHttpRequest. The data displayed on the UI by the API results is less critical than other data displayed on the same web page, so it is acceptable for some requests to not have the API data

displayed in the UI. However, calls made to the API should not delay rendering of other parts of the user

interface. You want your application to perform well when the API response is an error or a timeout.

What should you do?

A. Set the asynchronous option for your requests to the API to false and omit the widget displaying the API results when a timeout or error is encountered.

B. Set the asynchronous option for your request to the API to true and omit the widget displaying the API results when a timeout or error is encountered.

C. Catch timeout or error exceptions from the API call and keep trying with exponential backoff until the API response is successful.

D. Catch timeout or error exceptions from the API call and display the error response in the UI widget.

Correct Answer: A

QUESTION 3

Your team is building an application for a financial institution. The application\\'s frontend runs on Compute Engine, and the data resides in Cloud SQL and one Cloud Storage bucket. The application will collect data containing PII, which will be stored in the Cloud SQL database and the Cloud Storage bucket. You need to secure the PII data. What should you do?

A. 1) Create the relevant firewall rules to allow only the frontend to communicate with the Cloud SQL database 2) Using IAM, allow only the frontend service account to access the Cloud Storage bucket

B. 1) Create the relevant firewall rules to allow only the frontend to communicate with the Cloud SQL database 2) Enable private access to allow the frontend to access the Cloud Storage bucket privately

C. 1) Configure a private IP address for Cloud SQL 2) Use VPC-SC to create a service perimeter 3) Add the Cloud SQL database and the Cloud Storage bucket to the same service perimeter

D. 1) Configure a private IP address for Cloud SQL 2) Use VPC-SC to create a service perimeter 3) Add the Cloud SQL database and the Cloud Storage bucket to different service perimeters

Correct Answer: C

QUESTION 4

You have a container deployed on Google Kubernetes Engine. The container can sometimes be slow to launch, so you have implemented a liveness probe. You notice that the liveness probe occasionally fails on launch. What should you do?

A. Add a startup probe.

B. Increase the initial delay for the liveness probe.

C. Increase the CPU limit for the container.

D. Add a readiness probe.

Correct Answer: B

Explanation: https://kubernetes.io/docs/tasks/configure-pod-container/configure-liveness- readiness-startup-probes/#configure-probes

QUESTION 5

You deployed a new application to Google Kubernetes Engine and are experiencing some performance degradation. Your logs are being written to Cloud Logging, and you are using a Prometheus sidecar model for capturing metrics. You need to correlate the metrics and data from the logs to troubleshoot the performance issue and send real-time alerts while minimizing costs. What should you do?

A. Create custom metrics from the Cloud Logging logs, and use Prometheus to import the results using the Cloud Monitoring REST API.

B. Export the Cloud Logging logs and the Prometheus metrics to Cloud Bigtable. Run a query to join the results, and analyze in Google Data Studio.

C. Export the Cloud Logging logs and stream the Prometheus metrics to BigQuery. Run a recurring query to join the results, and send notifications using Cloud Tasks.

D. Export the Prometheus metrics and use Cloud Monitoring to view them as external metrics. Configure Cloud Monitoring to create log-based metrics from the logs, and correlate them with the Prometheus data.

Correct Answer: D

Reference: https://cloud.google.com/blog/products/operations/troubleshoot-gke-faster-with- monitoring-data-in-your-logs

QUESTION 6

You are a SaaS provider deploying dedicated blogging software to customers in your Google Kubernetes Engine (GKE) cluster. You want to configure a secure multi-tenant platform to ensure that each customer has access to only their own blog and can\\'t affect the workloads of other customers. What should you do?

A. Enable Application-layer Secrets on the GKE cluster to protect the cluster.

B. Deploy a namespace per tenant and use Network Policies in each blog deployment.

C. Use GKE Audit Logging to identify malicious containers and delete them on discovery.

D. Build a custom image of the blogging software and use Binary Authorization to prevent untrusted image deployments.

Correct Answer: B

Reference: https://cloud.google.com/kubernetes-engine/docs/concepts/multitenancy- overview

QUESTION 7

Your team develops stateless services that run on Google Kubernetes Engine (GKE). You need to deploy a new service that will only be accessed by other services running in the GKE cluster. The service will need to scale as quickly as possible to respond to changing load. What should you do?

A. Use a Vertical Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.

B. Use a Vertical Pod Autoscaler to scale the containers, and expose them via a NodePort Service.

C. Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.

D. Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a NodePort Service.

Correct Answer: C

Explanation: https://cloud.google.com/kubernetes-engine/docs/concepts/service

QUESTION 8

You are planning to add unit tests to your application. You need to be able to assert that published Pub/Sub messages are processed by your subscriber in order. You want the unit tests to be cost-effective and reliable. What should you do?

A. Implement a mocking framework.

B. Create a topic and subscription for each tester.

C. Add a filter by tester to the subscription.

D. Use the Pub/Sub emulator.

Correct Answer: D

Explanation: https://cloud.google.com/pubsub/docs/emulator, "Testing apps locally with the emulator".

QUESTION 9

You are planning to deploy your application in a Google Kubernetes Engine (GKE) cluster. Your application can scale horizontally, and each instance of your application needs to have a stable network identity and its own persistent disk.

Which GKE object should you use?

A. Deployment

B. StatefulSet

C. ReplicaSet

D. ReplicaController

Correct Answer: B

Reference: https://livebook.manning.com/book/kubernetes-in-action/chapter-10/46

QUESTION 10

You are running a containerized application on Google Kubernetes Engine. Your container images are stored in Container Registry. Your team uses CI/CD practices. You need to prevent the deployment of containers with known critical vulnerabilities. What should you do?

A. ?Use Web Security Scanner to automatically crawl your application ?Review your application logs for scan results, and provide an attestation that the container is free of known critical vulnerabilities ?Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed

B. ?Use Web Security Scanner to automatically crawl your application ?Review the scan results in the scan details page in the Cloud Console, and provide an attestation that the container is free of known critical vulnerabilities ?Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed

C. ?Enable the Container Scanning API to perform vulnerability scanning ?Review vulnerability reporting in Container Registry in the Cloud Console, and provide an attestation that the container is free of known critical vulnerabilities ?Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed

D. ?Enable the Container Scanning API to perform vulnerability scanning ?Programmatically review vulnerability reporting through the Container Scanning API, and provide an attestation that the container is free of known critical vulnerabilities ?Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed

Correct Answer: D

https://cloud.google.com/binary-authorization/docs/creating-attestations-kritis https://cloud.google.com/container-analysis/docs/os-overview

QUESTION 11

You are a developer at a large organization. You are deploying a web application to Google Kubernetes Engine (GKE). The DevOps team has built a CI/CD pipeline that uses Cloud Deploy to deploy the application to Dev Test, and Prod clusters in GKE. After Cloud Deploy successfully deploys the application to the Dev cluster you want to automatically promote it to the Test Cluster. How should you configure this process following Google-recommended best practices?

A. 1 Create a Cloud Build trigger that listens for SUCCEEDED Pub/Sub messages from the clouddeploy-operations topic. 2 Configure Cloud Build to include a step that promotes the application to the Test cluster

B. 1 Create a Cloud Function that calls the Google Cloud Deploy API to promote the application to the Test cluster 2 Configure this function to be triggered by SUCCEEDED Pub/Sub messages from the cloud-builds topic

C. 1 Create a Cloud Function that calls the Google Cloud Deploy API to promote the application to the Test cluster 2 Configure this function to be triggered by SUCCEEDED Pub/Sub messages from the clouddeploy operations topic

D. 1 Create a Cloud Build pipeline that uses the gke-deploy builder 2 Create a Cloud Build trigger that listens to SUCCEEDED Pub/Sub messages from the cloud-builds topic 3 Configure this pipeline to run a deployment step to the Test cluster

Correct Answer: B

QUESTION 12

You are building a CI/CD pipeline that consists of a version control system, Cloud Build, and Container Registry. Each time a new tag is pushed to the repository, a Cloud Build job is triggered, which runs unit tests on the new code builds a new Docker container image, and pushes it into Container Registry. The last step of your pipeline should deploy the new container to your production Google Kubernetes Engine (GKE) cluster. You need to select a tool and deployment strategy that meets the following requirements:

?Zero downtime is incurred ?Testing is fully automated ?Allows for testing before being rolled out to users ?Can quickly rollback if needed What should you do?

A. Trigger a Spinnaker pipeline configured as an A/B test of your new code and, if it is successful, deploy the container to production.

B. Trigger a Spinnaker pipeline configured as a canary test of your new code and, if it is successful, deploy the container to production.

C. Trigger another Cloud Build job that uses the Kubernetes CLI tools to deploy your new container to your GKE cluster, where you can perform a canary test.

D. Trigger another Cloud Build job that uses the Kubernetes CLI tools to deploy your new container to your GKE cluster, where you can perform a shadow test.

Correct Answer: D

Explanation: https://cloud.google.com/architecture/implementing-deployment-and-testing- strategies-on-gke#perform_a_shadow_test With a shadow test, you test the new version of your application by mirroring user traffic from the current application version without impacting the user requests.

QUESTION 13

You are designing a resource-sharing policy for applications used by different teams in a Google Kubernetes Engine cluster. You need to ensure that all applications can access the resources needed to run. What should you do? (Choose two.)

A. Specify the resource limits and requests in the object specifications.

B. Create a namespace for each team, and attach resource quotas to each namespace.

C. Create a LimitRange to specify the default compute resource requirements for each namespace.

D. Create a Kubernetes service account (KSA) for each application, and assign each KSA to the namespace.

E. Use the Anthos Policy Controller to enforce label annotations on all namespaces. Use taints and tolerations to allow resource sharing for namespaces.

Correct Answer: BC

Explanation: https://kubernetes.io/docs/concepts/policy/resource-quotas/ https://kubernetes.io/docs/concepts/policy/limit-range/ https://cloud.google.com/blog/products/containers-kubernetes/kubernetes-best-practices- resource-requests-andlimits

QUESTION 14

You need to deploy a new European version of a website hosted on Google Kubernetes Engine. The current and new websites must be accessed via the same HTTP(S) load balancer\\'s external IP address, but have different domain names. What should you do?

A. Define a new Ingress resource with a host rule matching the new domain

B. Modify the existing Ingress resource with a host rule matching the new domain

C. Create a new Service of type LoadBalancer specifying the existing IP address as the loadBalancerIP

D. Generate a new Ingress resource and specify the existing IP address as the kubernetes.io/ingress.global-static-ip-name annotation value

Correct Answer: B

https://kubernetes.io/docs/concepts/services-networking/ingress/#name-based-virtual- hosting Name-based virtual hosts support routing HTTP traffic to multiple host names at the same IP address.

QUESTION 15

You are developing an application that reads credit card data from a Pub/Sub subscription. You have written code and completed unit testing. You need to test the Pub/Sub integration before deploying to Google Cloud. What should you do?

A. Create a service to publish messages, and deploy the Pub/Sub emulator. Generate random content in the publishing service, and publish to the emulator.

B. Create a service to publish messages to your application. Collect the messages from Pub/Sub in production, and replay them through the publishing service.

C. Create a service to publish messages, and deploy the Pub/Sub emulator. Collect the messages from Pub/Sub in production, and publish them to the emulator.

D. Create a service to publish messages, and deploy the Pub/Sub emulator. Publish a standard set of testing messages from the publishing service to the emulator.

Correct Answer: D