



"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

[Applying Site Reliability Engineering (SRE) principles to a service]

You encountered a major service outage that affected all users of the service for multiple hours. After several hours of incident management, the service returned to normal, and user access was restored. You need to provide an incident summary to relevant stakeholders following the Site Reliability Engineering recommended practices. What should you do first?

- A. Call individual stakeholders to explain what happened.
- B. Develop a post-mortem to be distributed to stakeholders.
- C. Send the Incident State Document to all the stakeholders.
- D. Require the engineer responsible to write an apology email to all stakeholders.

Answer: B

Question: 2

[Applying Site Reliability Engineering (SRE) principles to a service]

You need to define SLOs for a high-traffic web application. Customers are currently happy with the application performance and availability. Based on current measurement, the 90th percentile of latency is 160 ms and the 95th percentile of latency is 300 ms over a 28-day window. What latency SLO should you publish?

- A. 90th percentile - 150 ms 95th percentile - 290 ms
- B. 90th percentile - 160 ms 95th percentile - 300 ms
- C. 90th percentile - 190 ms 95th percentile - 330 ms
- D. 90th percentile - 300 ms 95th percentile - 450 ms

Answer: B

Question: 3

[Building and implementing CI/CD pipelines for a service]

You are creating a CI/CD pipeline in Cloud Build to build an application container image. The application code is stored in GitHub. Your company requires that production image builds are only run against the main branch and that the change control team approves all pushes to the main branch. You want the image build to be as automated as possible. What should you do?

Choose 2 answers.

- A. Create a trigger on the Cloud Build job. Set the repository event setting to Pull request.
- B. Add the owners file to the Included files filter on the trigger.
- C. Create a trigger on the Cloud Build job. Set the repository event setting to Push to a branch.
- D. Configure a branch protection rule for the main branch on the repository.
- E. Enable the Approval option on the trigger.

Answer: C,D

Question: 4

[Applying Site Reliability Engineering (SRE) principles to a service]

You need to define Service Level Objectives (SLOs) for a high-traffic multi-region web application. Customers expect the application to always be available and have fast response times. Customers are currently happy with the application performance and availability. Based on current measurement, you observe that the 90th percentile of latency is 120ms and the 95th percentile of latency is 275ms over a 28-day window. What latency SLO would you recommend to the team to publish?

- A. 90th percentile – 100ms95th percentile – 250ms
- B. 90th percentile – 120ms95th percentile – 275ms
- C. 90th percentile – 150ms95th percentile – 300ms
- D. 90th percentile – 250ms95th percentile – 400ms

Answer: C

Question: 5

[Building and implementing service monitoring strategies]

Your company runs services by using multiple globally distributed Google Kubernetes Engine (GKE) clusters Your operations team has set up workload monitoring that uses Prometheus-based tooling for metrics alerts: and generating dashboards This setup does not provide a method to view metrics globally across all clusters You need to implement a scalable solution to support global Prometheus querying and minimize management overhead What should you do?

- A. Configure Prometheus cross-service federation for centralized data access
- B. Configure workload metrics within Cloud Operations for GKE
- C. Configure Prometheus hierarchical federation for centralized data access
- D. Configure Google Cloud Managed Service for Prometheus

Answer: D

Question: 6

[Building and implementing service monitoring strategies]

You want to share a Cloud Monitoring custom dashboard with a partner team What should you do?

- A. Provide the partner team with the dashboard URL to enable the partner team to create a copy of the dashboard
- B. Export the metrics to BigQuery Use Looker Studio to create a dashboard, and share the dashboard with the partner team
- C. Copy the Monitoring Query Language (MQL) query from the dashboard; and send the MQL query to the partner

team

D. Download the JSON definition of the dashboard, and send the JSON file to the partner team

Answer: D

Question: 7

[Building and implementing CI/CD pipelines for a service]

You have an application that runs in Google Kubernetes Engine (GKE). The application consists of several microservices that are deployed to GKE by using Deployments and Services. One of the microservices is experiencing an issue where a Pod returns 403 errors after the Pod has been running for more than five hours. Your development team is working on a solution but the issue will not be resolved for a month. You need to ensure continued operations until the microservice is fixed. You want to follow Google-recommended practices and use the fewest number of steps. What should you do?

- A. Create a cron job to terminate any Pods that have been running for more than five hours.
- B. Add a HTTP liveness probe to the microservice's deployment.
- C. Monitor the Pods and terminate any Pods that have been running for more than five hours.
- D. Configure an alert to notify you whenever a Pod returns 403 errors.

Answer: B

Question: 8

[Building and implementing CI/CD pipelines for a service]

You are building and deploying a microservice on Cloud Run for your organization. Your service is used by many applications internally. You are deploying a new release, and you need to test the new version extensively in the staging and production environments. You must minimize user and developer impact. What should you do?

- A. Deploy the new version of the service to the staging environment. Split the traffic, and allow 1 % of traffic through to the latest version. Test the latest version. If the test passes, gradually roll out the latest version to the staging and production environments.
- B. Deploy the new version of the service to the staging environment. Split the traffic, and allow 50% of traffic through to the latest version. Test the latest version. If the test passes, send all traffic to the latest version. Repeat for the production environment.
- C. Deploy the new version of the service to the staging environment with a new-release tag without serving traffic. Test the new-release version. If the test passes; gradually roll out this tagged version. Repeat for the production environment.
- D. Deploy a new environment with the green tag to use as the staging environment. Deploy the new version of the service to the green environment and test the new version. If the tests pass, send all traffic to the green environment and delete the existing staging environment. Repeat for the production environment.

Answer: C

Question: 9

[Optimizing service performance]

Your company processes IOT data at scale by using Pub/Sub, App Engine standard environment, and an application written in GO. You noticed that the performance inconsistently degrades at peak load. You could not reproduce this issue on your workstation. You need to continuously monitor the application in production to identify slow paths in the code. You want to minimize performance impact and management overhead. What should you do?

- A. Install a continuous profiling tool into Compute Engine. Configure the application to send profiling data to the tool.
- B. Periodically run the go tool pprof command against the application instance. Analyze the results by using flame graphs.
- C. Configure Cloud Profiler, and initialize the cloud.go@gle.com/go/profiler library in the application.
- D. Use Cloud Monitoring to assess the App Engine CPU utilization metric.

Answer: C

Question: 10

[Building and implementing CI/CD pipelines for a service]

Your company uses a CI/CD pipeline with Cloud Build and Artifact Registry to deploy container images to Google Kubernetes Engine (GKE). Images are tagged with the latest commit hash and promoted to production after successful testing in the development and pre-production environments. A recent production deployment caused the application to fail due to untested integration functionality, requiring a disruptive manual rollback. During the rollback, you noticed many old and unused container images accumulating in Artifact Registry. You need to improve rollout and rollback management and clean up the old container images. What should you do?

- A. Adopt Cloud Deploy for managing deployments, and schedule a Cloud Build job for container image cleanup.
- B. Deploy Cloud Service Mesh across the GKE clusters, and manually clean up Artifact Registry images.
- C. Adopt Cloud Deploy for managing deployments, and implement an Artifact Registry cleanup policy.
- D. Set up a rollback pipeline in Cloud Build, and implement an Artifact Registry cleanup policy.

Answer: C

Question: 11

[Applying Site Reliability Engineering (SRE) principles to a service]

You work for a global organization and run a service with an availability target of 99% with limited engineering resources. For the current calendar month you noticed that the service has 99.5% availability. You must ensure that your service

meets the defined availability goals and can react to business changes including the upcoming launch of new features You also need to reduce technical debt while minimizing operational costs You want to follow Google-recommended practices What should you do?

- A. Add N+1 redundancy to your service by adding additional compute resources to the service
- B. Identify, measure and eliminate toil by automating repetitive tasks
- C. Define an error budget for your service level availability and minimize the remaining error budget
- D. Allocate available engineers to the feature backlog while you ensure that the service remains within the availability target

Answer: B

Question: 12

[Building and implementing service monitoring strategies]

You currently store the virtual machine (VM) utilization logs in Stackdriver. You need to provide an easy-to-share interactive VM utilization dashboard that is updated in real time and contains information aggregated on a quarterly basis. You want to use Google Cloud Platform solutions. What should you do?

- A. 1. Export VM utilization logs from Stackdriver to BigQuery.2. Create a dashboard in Data Studio.3. Share the dashboard with your stakeholders.
- B. 1. Export VM utilization logs from Stackdriver to Cloud Pub/Sub.2. From Cloud Pub/Sub, send the logs to a Security Information and Event Management (SIEM) system.3. Build the dashboards in the SIEM system and share with your stakeholders.
- C. 1. Export VM utilization logs from Stackdriver to BigQuery.2. From BigQuery, export the logs to a CSV file.3. Import the CSV file into Google Sheets.4. Build a dashboard in Google Sheets and share it with your stakeholders.
- D. 1. Export VM utilization logs from Stackdriver to a Cloud Storage bucket.2. Enable the Cloud Storage API to pull the logs programmatically.3. Build a custom data visualization application.4. Display the pulled logs in a custom dashboard.

Answer: A

Question: 13

[Building and implementing service monitoring strategies]

Your company is creating a new cloud-native Google Cloud organization. You expect this Google Cloud organization to first be used by a small number of departments and then expand to be used by a large number of departments. Each department has a large number of applications varying in size. You need to design the VPC network architecture. Your solution must minimize the amount of management required while remaining flexible enough for development teams to quickly adapt to their evolving needs. What should you do?

- A. Create a separate VPC for each department and connect the VPCs with VPC Network Peering.
- B. Create a separate VPC for each department and use Private Service Connect to connect the VPCs.
- C. Create a separate VPC for each application and use Private Service Connect to connect the VPCs.
- D. Create a separate VPC for each department and connect the VPCs with Cloud VPN.

Answer: A

Question: 14

[Building and implementing service monitoring strategies]

Your company operates in a highly regulated domain that requires you to store all organization logs for seven years. You want to minimize logging infrastructure complexity by using managed services. You need to avoid any future loss of log capture or stored logs due to misconfiguration or human error. What should you do?

- A. Use Cloud Logging to configure an aggregated sink at the organization level to export all logs into a BigQuery dataset
- B. Use Cloud Logging to configure an aggregated sink at the organization level to export all logs into Cloud Storage with a seven-year retention policy and Bucket Lock
- C. Use Cloud Logging to configure an export sink at each project level to export all logs into a BigQuery dataset
- D. Use Cloud Logging to configure an export sink at each project level to export all logs into Cloud Storage with a seven-year retention policy and Bucket Lock

Answer: B

Question: 15

[Building and implementing CI/CD pipelines for a service]

You are configuring the frontend tier of an application deployed in Google Cloud. The frontend tier is hosted in `ngmx` and deployed using a managed instance group with an Envoy-based external HTTP(S) load balancer in front. The application is deployed entirely within the `eu-west-2` region and only serves users based in the United Kingdom. You need to choose the most cost-effective network tier and load balancing configuration. What should you use?

- A. Premium Tier with a global load balancer
- B. Premium Tier with a regional load balancer
- C. Standard Tier with a global load balancer
- D. Standard Tier with a regional load balancer

Answer: B

Question: 16

[Applying Site Reliability Engineering (SRE) principles to a service]

You support a web application that is hosted on Compute Engine. The application provides a booking service for thousands of users. Shortly after the release of a new feature, your monitoring dashboard shows that all users are experiencing latency at login. You want to mitigate the impact of the incident on the users of your service. What should you do first?

- A. Roll back the recent release.
- B. Review the Stackdriver monitoring.
- C. Upsize the virtual machines running the login services.
- D. Deploy a new release to see whether it fixes the problem.

Answer: C

Question: 17

[Building and implementing CI/CD pipelines for a service]

You are configuring a CI pipeline. The build step for your CI pipeline integration testing requires access to APIs inside your private VPC network. Your security team requires that you do not expose API traffic publicly. You need to implement a solution that minimizes management overhead. What should you do?

- A. Use Cloud Build private pools to connect to the private VPC.
- B. Use Spinnaker for Google Cloud to connect to the private VPC.
- C. Use Cloud Build as a pipeline runner. Configure Internal HTTP(S) Load Balancing for API access.
- D. Use Cloud Build as a pipeline runner. Configure External HTTP(S) Load Balancing with a Google Cloud Armor policy for API access.

Answer: A

Question: 18

[Building and implementing CI/CD pipelines for a service]

You are running a web application deployed to a Compute Engine managed instance group. Ops Agent is installed on all instances. You recently noticed suspicious activity from a specific IP address. You need to configure Cloud Monitoring to view the number of requests from that specific IP address with minimal operational overhead. What should you do?

- A. Configure the Ops Agent with a logging receiver. Create a logs-based metric.
- B. Create a script to scrape the web server log. Export the IP address request metrics to the Cloud Monitoring API.
- C. Update the application to export the IP address request metrics to the Cloud Monitoring API.
- D. Configure the Ops Agent with a metrics receiver.

Answer: A

Question: 19

[Building and implementing CI/CD pipelines for a service]

Your organization is starting to containerize with Google Cloud. You need a fully managed storage solution for container

images and Helm charts. You need to identify a storage solution that has native integration into existing Google Cloud services, including Google Kubernetes Engine (GKE), Cloud Run, VPC Service Controls, and Identity and Access Management (IAM). What should you do?

- A. Use Docker to configure a Cloud Storage driver pointed at the bucket owned by your organization.
- B. Configure Container Registry as an OCI-based container registry for container images.
- C. Configure Artifact Registry as an OCI-based container registry for both Helm charts and container images.
- D. Configure an open source container registry server to run in GKE with a restrictive role-based access control (RBAC) configuration.

Answer: C

Question: 20

[Building and implementing CI/CD pipelines for a service]

You are developing reusable infrastructure as code modules. Each module contains integration tests that launch the module in a test project. You are using GitHub for source control. You need to continuously test your feature branch and ensure that all code is tested before changes are accepted. You need to implement a solution to automate the integration tests. What should you do?

- A. Use a Jenkins server for CI/CD pipelines. Periodically run all tests in the feature branch.
- B. Use Cloud Build to run the tests. Trigger all tests to run after a pull request is merged.
- C. Ask the pull request reviewers to run the integration tests before approving the code.
- D. Use Cloud Build to run tests in a specific folder. Trigger Cloud Build for every GitHub pull request.

Answer: D

Question: 21

[Building and implementing service monitoring strategies]

You have a pool of application servers running on Compute Engine. You need to provide a secure solution that requires the least amount of configuration and allows developers to easily access application logs for troubleshooting. How would you implement the solution on GCP?

- A. • Deploy the Stackdriver logging agent to the application servers. • Give the developers the IAM Logs Viewer role to access Stackdriver and view logs.
- B. • Deploy the Stackdriver logging agent to the application servers. • Give the developers the IAM Logs Private Logs Viewer role to access Stackdriver and view logs.
- C. • Deploy the Stackdriver monitoring agent to the application servers. • Give the developers the IAM Monitoring Viewer role to access Stackdriver and view metrics.

D. • Install the gsutil command line tool on your application servers. • Write a script using gsutil to upload your application log to a Cloud Storage bucket, and then schedule it to run via cron every 5 minutes. • Give the developers IAM Object Viewer access to view the logs in the specified bucket.

Answer: A

Question: 22

[Building and implementing CI/CD pipelines for a service]

Your application artifacts are being built and deployed via a CI/CD pipeline. You want the CI/CD pipeline to securely access application secrets. You also want to more easily rotate secrets in case of a security breach. What should you do?

- A. Prompt developers for secrets at build time. Instruct developers to not store secrets at rest.
- B. Store secrets in a separate configuration file on Git. Provide select developers with access to the configuration file.
- C. Store secrets in Cloud Storage encrypted with a key from Cloud KMS. Provide the CI/CD pipeline with access to Cloud KMS via IAM.
- D. Encrypt the secrets and store them in the source code repository. Store a decryption key in a separate repository and grant your pipeline access to it

Answer: C

Question: 23

[Building and implementing CI/CD pipelines for a service]

Your application images are built using Cloud Build and pushed to Google Container Registry (GCR). You want to be able to specify a particular version of your application for deployment based on the release version tagged in source control. What should you do when you push the image?

- A. Reference the image digest in the source control tag.
- B. Supply the source control tag as a parameter within the image name.
- C. Use Cloud Build to include the release version tag in the application image.
- D. Use GCR digest versioning to match the image to the tag in source control.

Answer: B

Question: 24

[Building and implementing CI/CD pipelines for a service]

You are the on-call Site Reliability Engineer for a microservice that is deployed to a Google Kubernetes Engine (GKE) Autopilot cluster. Your company runs an online store that publishes order messages to Pub/Sub and a microservice receives these messages and updates stock information in the warehousing system. A sales event caused an increase in orders, and the stock information is not being updated quickly enough. This is causing a large number of orders to be accepted for

products that are out of stock You check the metrics for the microservice and compare them to typical levels.

Microservice metrics	Typical state	Current state
Average CPU across all Pods	20% of Pod limit	30% of Pod limit
Average memory across all Pods	10% of Pod limit	10% of Pod limit
Pub Sub subscription Average oldest unacknowledged message age	347 milliseconds	8074 milliseconds
Pub, Sub subscription Average undelivered messages	5 messages	14705 messages
Pub Sub subscription Average acknowledgment latency	312 milliseconds	354 milliseconds

You need to ensure that the warehouse system accurately reflects product inventory at the time orders are placed and minimize the impact on customers What should you do?

- A. Decrease the acknowledgment deadline on the subscription
- B. Add a virtual queue to the online store that allows typical traffic levels
- C. Increase the number of Pod replicas
- D. Increase the Pod CPU and memory limits

Answer: C

Question: 25

[Building and implementing service monitoring strategies]

You support a service that recently had an outage. The outage was caused by a new release that exhausted the service memory resources. You rolled back the release successfully to mitigate the impact on users. You are now in charge of the post-mortem for the outage. You want to follow Site Reliability Engineering practices when developing the post-mortem. What should you do?

- A. Focus on developing new features rather than avoiding the outages from recurring.
- B. Focus on identifying the contributing causes of the incident rather than the individual responsible for the cause.
- C. Plan individual meetings with all the engineers involved. Determine who approved and pushed the new release to production.
- D. Use the Git history to find the related code commit. Prevent the engineer who made that commit from working on production services.

Answer: B

Question: 26

[Building and implementing service monitoring strategies]

You support a production service that runs on a single Compute Engine instance. You regularly need to spend time on recreating the service by deleting the crashing instance and creating a new instance based on the relevant image. You want to reduce the time spent performing manual operations while following Site Reliability Engineering principles.

What should you do?

- A. File a bug with the development team so they can find the root cause of the crashing instance.
- B. Create a Managed Instance Group with a single instance and use health checks to determine the system status.
- C. Add a Load Balancer in front of the Compute Engine instance and use health checks to determine the system status.
- D. Create a Stackdriver Monitoring dashboard with SMS alerts to be able to start recreating the crashed instance promptly after it has crashed.

Answer: B

Question: 27

[Applying Site Reliability Engineering (SRE) principles to a service]

You support a web application that runs on App Engine and uses CloudSQL and Cloud Storage for data storage. After a short spike in website traffic, you notice a big increase in latency for all user requests, increase in CPU use, and the number of processes running the application. Initial troubleshooting reveals:

After the initial spike in traffic, load levels returned to normal but users still experience high latency. Requests for content from the CloudSQL database and images from Cloud Storage show the same high latency.

No changes were made to the website around the time the latency increased.

There is no increase in the number of errors to the users.

You expect another spike in website traffic in the coming days and want to make sure users don't experience latency. What should you do?

- A. Upgrade the GCS buckets to Multi-Regional.
- B. Enable high availability on the CloudSQL instances.
- C. Move the application from App Engine to Compute Engine.
- D. Modify the App Engine configuration to have additional idle instances.

Answer: D

Question: 28

[Applying Site Reliability Engineering (SRE) principles to a service]

You support a user-facing web application. When analyzing the application's error budget over the previous six months, you notice that the application has never consumed more than 5% of its error budget in any given time window. You hold a Service Level Objective (SLO) review with business stakeholders and confirm that the SLO is set appropriately. You want your application's SLO to more closely reflect its observed reliability. What steps can you take to further that goal while balancing velocity, reliability, and business needs? (Choose two.)

- A. Add more serving capacity to all of your application's zones.
- B. Have more frequent or potentially risky application releases.
- C. Tighten the SLO match the application's observed reliability.
- D. Implement and measure additional Service Level Indicators (SLIs) from the application.
- E. Announce planned downtime to consume more error budget, and ensure that users are not depending on a tighter SLO.

Answer: D,E

Question: 29

[Building and implementing service monitoring strategies]

Your organization wants to collect system logs that will be used to generate dashboards in Cloud Operations for their Google Cloud project. You need to configure all current and future Compute Engine instances to collect the system logs and you must ensure that the Ops Agent remains up to date. What should you do?

- A. Use the gcloud CLI to install the Ops Agent on each VM listed in the Cloud Asset Inventory
- B. Select all VMs with an Agent status of Not detected on the Cloud Operations VMs dashboard Then select Install agents
- C. Use the gcloud CLI to create an Agent Policy.
- D. Install the Ops Agent on the Compute Engine image by using a startup script

Answer: C

Question: 30

[Applying Site Reliability Engineering (SRE) principles to a service]

Your organization wants to increase the availability target of an application from 99.9% to 99.99% for an investment of \$2,000. The application's current revenue is \$1,000,000. You need to determine whether the increase in availability is worth the investment for a single year of usage. What should you do?

- A. Calculate the value of improved availability to be \$900, and determine that the increase in availability is not worth the investment
- B. Calculate the value of improved availability to be \$1,000 and determine that the increase in availability is not worth the investment
- C. Calculate the value of improved availability to be \$1,000 and determine that the increase in availability is worth the investment
- D. Calculate the value of improved availability to be \$9,000. and determine that the increase in availability is worth the investment

Answer: A

Question: 31

[Optimizing service performance]

Your organization stores all application logs from multiple Google Cloud projects in a central Cloud Logging project. Your security team wants to enforce a rule that each project team can only view their respective logs, and only the operations team can view all the logs. You need to design a solution that meets the security team's requirements, while minimizing costs. What should you do?

- A. Export logs to BigQuery tables for each project team. Grant project teams access to their tables. Grant logs writer access to the operations team in the central logging project.
- B. Create log views for each project team, and only show each project team their application logs. Grant the operations team access to the _ All-logs View in the central logging project.
- C. Grant each project team access to the project _ Default view in the central logging project. Grant logging viewer access to the operations team in the central logging project.
- D. Create Identity and Access Management (IAM) roles for each project team and restrict access to the _ Default log view in their individual Google Cloud project. Grant viewer access to the operations team in the central logging project.

Answer: B

Question: 32

[Applying Site Reliability Engineering (SRE) principles to a service]

Your company is migrating its production systems to Google Cloud. You need to implement site reliability engineering (SRE) practices during the migration to minimize customer impact from potential future incidents. Which two SRE practices should you implement?

Choose 2 answers

- A. Ensure that full autonomy and permissions are only granted to the on-call team.
- B. Automate common tasks to analyze key impact information and intelligently suggest mitigating actions for the on-call team.
- C. Ensure that all teams can modify the production environment to resolve issues.
- D. Create an alerting mechanism for your SRE team based on your system's internal behavior.
- E. Create up-to-date playbooks with instructions for debugging and mitigating issues.

Answer: B,E

Question: 33

[Building and implementing CI/CD pipelines for a service]

Your team is designing a new application for deployment both inside and outside Google Cloud Platform (GCP). You need to collect detailed metrics such as system resource utilization. You want to use centralized GCP services while minimizing the amount of work required to set up this collection system. What should you do?

- A. Import the Stackdriver Profiler package, and configure it to relay function timing data to Stackdriver for further analysis.
- B. Import the Stackdriver Debugger package, and configure the application to emit debug messages with timing information.
- C. Instrument the code using a timing library, and publish the metrics via a health check endpoint that is scraped by Stackdriver.
- D. Install an Application Performance Monitoring (APM) tool in both locations, and configure an export to a central data storage location for analysis.

Answer: A

Question: 34

[Building and implementing CI/CD pipelines for a service]

You are investigating issues in your production application that runs on Google Kubernetes Engine (GKE). You determined that the source of the issue is a recently updated container image, although the exact change in code was not identified. The deployment is currently pointing to the latest tag. You need to update your cluster to run a version of the container that functions as intended. What should you do?

- A. Create a new tag called stable that points to the previously working container, and change the deployment to point to the new tag.
- B. Apply the latest tag to the previous container image, and do a rolling update on the deployment.
- C. Build a new container from a previous Git tag, and do a rolling update on the deployment to the new container.
- D. Alter the deployment to point to the sha256 digest of the previously working container.

Answer: D

Question: 35

[Applying Site Reliability Engineering (SRE) principles to a service]

You are responsible for the reliability of a high-volume enterprise application. A large number of users report that an important subset of the application's functionality – a data intensive reporting feature – is consistently failing with an HTTP 500 error. When you investigate your application's dashboards, you notice a strong correlation between the failures and a metric that represents the size of an internal queue used for generating reports. You trace the failures to a reporting backend that is experiencing high I/O wait times. You quickly fix the issue by resizing the backend's persistent disk (PD). How would you need to create an availability Service Level Indicator (SLI) for the report generation feature. How would you define it?

- A. As the I/O wait times aggregated across all report generation backends
- B. As the proportion of report generation requests that result in a successful response
- C. As the application's report generation queue size compared to a known-good threshold
- D. As the reporting backend PD throughput capacity compared to a known-good threshold

Answer: B

Question: 36

[Building and implementing CI/CD pipelines for a service]

You are reviewing your deployment pipeline in Google Cloud Deploy. You must reduce toil in the pipeline and you want to minimize the amount of time it takes to complete an end-to-end deployment. What should you do?

Choose 2 answers

- A. Create a trigger to notify the required team to complete the next step when manual intervention is required
- B. Divide the automation steps into smaller tasks
- C. Use a script to automate the creation of the deployment pipeline in Google Cloud Deploy
- D. Add more engineers to finish the manual steps.
- E. Automate promotion approvals from the development environment to the test environment

Answer: A,E

Question: 37

[Building and implementing service monitoring strategies]

You support an application running on App Engine. The application is used globally and accessed from various device types.

You want to know the number of connections. You are using Stackdriver Monitoring for App Engine. What metric should you use?

- A. flex/connections/current
- B. tcp_ssl_proxy/new_connections
- C. tcp_ssl_proxy/open_connections
- D. flex/instance/connections/current

Answer: A

Question: 38

[Applying Site Reliability Engineering (SRE) principles to a service]

You are part of an organization that follows SRE practices and principles. You are taking over the management of a new service from the Development Team, and you conduct a Production Readiness Review (PRR). After the PRR analysis phase, you determine that the service cannot currently meet its Service Level Objectives (SLOs). You want to ensure that the service can meet its SLOs in production. What should you do next?

- A. Adjust the SLO targets to be achievable by the service so you can bring it into production.
- B. Notify the development team that they will have to provide production support for the service.
- C. Identify recommended reliability improvements to the service to be completed before handover.
- D. Bring the service into production with no SLOs and build them when you have collected operational data.

Answer: C

Question: 39

[Applying Site Reliability Engineering (SRE) principles to a service]

Your team is writing a postmortem after an incident on your external facing application Your team wants to improve the

postmortem policy to include triggers that indicate whether an incident requires a postmortem Based on Site Reliability Engineering (SRE) practices, what triggers should be defined in the postmortem policy?
Choose 2 answers

- A. An external stakeholder asks for a postmortem
- B. Data is lost due to an incident
- C. An internal stakeholder requests a postmortem
- D. The monitoring system detects that one of the instances for your application has failed
- E. The CD pipeline detects an issue and rolls back a problematic release.

Answer: A,C

Question: 40

[Applying Site Reliability Engineering (SRE) principles to a service]

You recently noticed that one of your services has exceeded the error budget for the current rolling window period. Your company's product team is about to launch a new feature. You want to follow Site Reliability Engineering (SRE) practices. What should you do?

- A. Notify the team about the lack of error budget and ensure that all their tests are successful so the launch will not further risk the error budget.
- B. Notify the team that their error budget is used up. Negotiate with the team for a launch freeze or tolerate a slightly worse user experience.
- C. Escalate the situation and request additional error budget.
- D. Look through other metrics related to the product and find SLOs with remaining error budget. Reallocate the error budgets and allow the feature launch.

Answer: B

Question: 41

[Building and implementing CI/CD pipelines for a service]

You are deploying an application that needs to access sensitive information. You need to ensure that this information is encrypted and the risk of exposure is minimal if a breach occurs. What should you do?

- A. Store the encryption keys in Cloud Key Management Service (KMS) and rotate the keys frequently
- B. Inject the secret at the time of instance creation via an encrypted configuration management system.
- C. Integrate the application with a Single sign-on (SSO) system and do not expose secrets to the application
- D. Leverage a continuous build pipeline that produces multiple versions of the secret for each instance of the application.

Answer: A

Question: 42

[Building and implementing CI/CD pipelines for a service]

The new version of your containerized application has been tested and is ready to be deployed to production on Google Kubernetes Engine (GKE). You could not fully load-test the new version in your pre-production environment and you need to ensure that the application does not have performance problems after deployment. Your deployment must be automated. What should you do?

- A. Deploy the application through a continuous delivery pipeline by using canary deployments. Use Cloud Monitoring to look for performance issues, and ramp up traffic as supported by the metrics.
- B. Deploy the application through a continuous delivery pipeline by using blue/green deployments. Migrate traffic to the new version of the application and use Cloud Monitoring to look for performance issues.
- C. Deploy the application by using `kubectl` and use Config Connector to slowly ramp up traffic between versions. Use Cloud Monitoring to look for performance issues.
- D. Deploy the application by using `kubectl` and set the `spec.updateStrategy.type` field to `RollingUpdate`. Use Cloud Monitoring to look for performance issues, and run the `kubectl rollback` command if there are any issues.

Answer: A

Question: 43

[Building and implementing service monitoring strategies]

You are designing a new Google Cloud organization for a client. Your client is concerned with the risks associated with long-lived credentials created in Google Cloud. You need to design a solution to completely eliminate the risks associated with the use of JSON service account keys while minimizing operational overhead. What should you do?

- A. Use custom versions of predefined roles to exclude all `iam.serviceAccountKeys.*` service account role permissions.
- B. Apply the constraints/`iam.disableServiceAccountKeycreation` constraint to the organization.
- C. Apply the constraints/`iam.disableServiceAccountKeyUp10ad` constraint to the organization.
- D. Grant the roles/`iam.serviceAccountKeyAdmin` IAM role to organization administrators only.

Answer: B

Question: 44

[Optimizing service performance]

You are the Site Reliability Engineer responsible for managing your company's data services and products. You regularly navigate operational challenges, such as unpredictable data volume and high cost, with your company's data ingestion

processes. You recently learned that a new data ingestion product will be developed in Google Cloud. You need to collaborate with the product development team to provide operational input on the new product. What should you do?

- A. Deploy the prototype product in a test environment, run a load test, and share the results with the product development team.
- B. When the initial product version passes the quality assurance phase and compliance assessments, deploy the product to a staging environment. Share error logs and performance metrics with the product development team.
- C. When the new product is used by at least one internal customer in production, share error logs and monitoring metrics with the product development team.
- D. Review the design of the product with the product development team to provide feedback early in the design phase.

Answer: D

Question: 45

[Building and implementing CI/CD pipelines for a service]

You use Spinnaker to deploy your application and have created a canary deployment stage in the pipeline. Your application has an in-memory cache that loads objects at start time. You want to automate the comparison of the canary version against the production version. How should you configure the canary analysis?

- A. Compare the canary with a new deployment of the current production version.
- B. Compare the canary with a new deployment of the previous production version.
- C. Compare the canary with the existing deployment of the current production version.
- D. Compare the canary with the average performance of a sliding window of previous production versions.

Answer: A

Question: 46

[Building and implementing CI/CD pipelines for a service]

You need to enforce several constraint templates across your Google Kubernetes Engine (GKE) clusters. The constraints include policy parameters, such as restricting the Kubernetes API. You must ensure that the policy parameters are stored in a GitHub repository and automatically applied when changes occur. What should you do?

- A. Set up a GitHub action to trigger Cloud Build when there is a parameter change. In Cloud Build, run a `gcloud` CLI command to apply the change.
- B. When there is a change in GitHub, use a web hook to send a request to Anthos Service Mesh, and apply the change.
- C. Configure Anthos Config Management with the GitHub repository. When there is a change in the repository, use Anthos Config Management to apply the change.
- D. Configure Config Connector with the GitHub repository. When there is a change in the repository, use Config

Connector to apply the change.

Answer: C

Question: 47

[Applying Site Reliability Engineering (SRE) principles to a service]

You support a multi-region web service running on Google Kubernetes Engine (GKE) behind a Global HTTP's Cloud Load Balancer (CLB). For legacy reasons, user requests first go through a third-party Content Delivery Network (CDN). which then routes traffic to the CLB. You have already implemented an availability Service Level Indicator (SLI) at the CLB level. However, you want to increase coverage in case of a potential load balancer misconfiguration. CDN failure, or other global networking catastrophe. Where should you measure this new SLI?

Choose 2 answers

- A. Your application servers' logs
- B. Instrumentation coded directly in the client
- C. Metrics exported from the application servers
- D. GKE health checks for your application servers
- E. A synthetic client that periodically sends simulated user requests

Answer: B,E

Question: 48

[Building and implementing CI/CD pipelines for a service]

You are building the CI/CD pipeline for an application deployed to Google Kubernetes Engine (GKE) The application is deployed by using a Kubernetes Deployment, Service, and Ingress The application team asked you to deploy the application by using the blue'green deployment methodology You need to implement the rollback actions What should you do?

- A. Run the kubectl rollout undo command
- B. Delete the new container image, and delete the running Pods
- C. Update the Kubernetes Service to point to the previous Kubernetes Deployment
- D. Scale the new Kubernetes Deployment to zero

Answer: C

Question: 49

[Building and implementing service monitoring strategies]

You support an application that stores product information in cached memory. For every cache miss, an entry is logged in

Stackdriver Logging. You want to visualize how often a cache miss happens over time. What should you do?

- A. Link Stackdriver Logging as a source in Google Data Studio. Filter (he logs on the cache misses).
- B. Configure Stackdriver Profiler to identify and visualize when the cache misses occur based on the logs.
- C. Create a logs-based metric in Stackdriver Logging and a dashboard for that metric in Stackdriver Monitoring.
- D. Configure BigQuery as a sink for Stackdriver Logging. Create a scheduled query to filter the cache miss logs and write them to a separate table

Answer: C

Question: 50

[Building and implementing service monitoring strategies]

You are currently planning how to display Cloud Monitoring metrics for your organization's Google Cloud projects. Your organization has three folders and six projects:

Folders	Projects
Development	• app-one-dev • app-two-dev
Staging	• app-one-staging • app-two-staging
Production	• app-one-prod • app-two-prod

You want to configure Cloud Monitoring dashboards to only display metrics from the projects within one folder. You need to ensure that the dashboards do not display metrics from projects in the other folders. You want to follow Google-recommended practices. What should you do?

- A. Create a single new scoping project
- B. Create new scoping projects for each folder
- C. Use the current app-one-prod project as the scoping project
- D. Use the current app-one-dev, app-one-staging and app-one-prod projects as the scoping project for each folder

Answer: A

Question: 51

[Applying Site Reliability Engineering (SRE) principles to a service]

You are writing a postmortem for an incident that severely affected users. You want to prevent similar incidents in the future. Which two of the following sections should you include in the postmortem? (Choose two.)

- A. An explanation of the root cause of the incident
- B. A list of employees responsible for causing the incident
- C. A list of action items to prevent a recurrence of the incident
- D. Your opinion of the incident's severity compared to past incidents
- E. Copies of the design documents for all the services impacted by the incident

Answer: A,C

Question: 52

[Managing service incidents]

You encounter a large number of outages in the production systems you support. You receive alerts for all the outages that wake you up at night. The alerts are due to unhealthy systems that are automatically restarted within a minute. You want to set up a process that would prevent staff burnout while following Site Reliability Engineering practices. What should you do?

- A. Eliminate unactionable alerts.
- B. Create an incident report for each of the alerts.
- C. Distribute the alerts to engineers in different time zones.
- D. Redefine the related Service Level Objective so that the error budget is not exhausted.

Answer: A

Question: 53

[Building and implementing service monitoring strategies]

Your Cloud Run application writes unstructured logs as text strings to Cloud Logging. You want to convert the unstructured logs to JSON-based structured logs. What should you do?

- A. Install a Fluent Bit sidecar container, and use a JSON parser.
- B. Install the log agent in the Cloud Run container image, and use the log agent to forward logs to Cloud Logging.
- C. Configure the log agent to convert log text payload to JSON payload.

D. Modify the application to use Cloud Logging software development kit (SDK), and send log entries with a jsonPayload field.

Answer: D

Question: 54

[Building and implementing service monitoring strategies]

You are running an application on Compute Engine and collecting logs through Stackdriver. You discover that some personally identifiable information (PII) is leaking into certain log entry fields. All PII entries begin with the text userinfo. You want to capture these log entries in a secure location for later review and prevent them from leaking to Stackdriver Logging. What should you do?

- A. Create a basic log filter matching userinfo, and then configure a log export in the Stackdriver console with Cloud Storage as a sink.
- B. Use a Fluentd filter plugin with the Stackdriver Agent to remove log entries containing userinfo, and then copy the entries to a Cloud Storage bucket.
- C. Create an advanced log filter matching userinfo, configure a log export in the Stackdriver console with Cloud Storage as a sink, and then configure a log exclusion with userinfo as a filter.
- D. Use a Fluentd filter plugin with the Stackdriver Agent to remove log entries containing userinfo, create an advanced log filter matching userinfo, and then configure a log export in the Stackdriver console with Cloud Storage as a sink.

Answer: B

Question: 55

[Building and implementing CI/CD pipelines for a service]

You recently deployed your application in Google Kubernetes Engine (GKE) and now need to release a new version of the application. You need the ability to instantly roll back to the previous version of the application in case there are issues with the new version. Which deployment model should you use?

- A. Perform a rolling deployment and test your new application after the deployment is complete.
- B. Perform A/B testing, and test your application periodically after the deployment is complete.
- C. Perform a canary deployment, and test your new application periodically after the new version is deployed.
- D. Perform a blue/green deployment and test your new application after the deployment is complete.

Answer: D

Question: 56

[Building and implementing CI/CD pipelines for a service]

You support an e-commerce application that runs on a large Google Kubernetes Engine (GKE) cluster deployed on-premises.

and on Google Cloud Platform. The application consists of microservices that run in containers. You want to identify containers that are using the most CPU and memory. What should you do?

- A. Use Stackdriver Kubernetes Engine Monitoring.
- B. Use Prometheus to collect and aggregate logs per container, and then analyze the results in Grafana.
- C. Use the Stackdriver Monitoring API to create custom metrics, and then organize your containers using groups.
- D. Use Stackdriver Logging to export application logs to BigQuery, aggregate logs per container, and then analyze CPU and memory consumption.

Answer: A

Question: 57

[Building and implementing service monitoring strategies]

You have migrated an e-commerce application to Google Cloud Platform (GCP). You want to prepare the application for the upcoming busy season. What should you do first to prepare for the busy season?

- A. Load test the application to profile its performance for scaling.
- B. Enable AutoScaling on the production clusters, in case there is growth.
- C. Pre-provision double the compute power used last season, expecting growth.
- D. Create a runbook on inflating the disaster recovery (DR) environment if there is growth.

Answer: A

Question: 58

[Building and implementing CI/CD pipelines for a service]

You need to deploy a new service to production. The service needs to automatically scale using a Managed Instance Group (MIG) and should be deployed over multiple regions. The service needs a large number of resources for each instance and you need to plan for capacity. What should you do?

- A. Use the n2-highcpu-96 machine type in the configuration of the MIG.
- B. Monitor results of Stackdriver Trace to determine the required amount of resources.
- C. Validate that the resource requirements are within the available quota limits of each region.
- D. Deploy the service in one region and use a global load balancer to route traffic to this region.

Answer: C

Question: 59

[Building and implementing CI/CD pipelines for a service]

You deploy a new release of an internal application during a weekend maintenance window when there is minimal user traffic. After the window ends, you learn that one of the new features isn't working as expected in the production

environment. After an extended outage, you roll back the new release and deploy a fix. You want to modify your release process to reduce the mean time to recovery so you can avoid extended outages in the future. What should you do?

Choose 2 answers

- A. Before merging new code, require 2 different peers to review the code changes.
- B. Adopt the blue/green deployment strategy when releasing new code via a CD server.
- C. Integrate a code linting tool to validate coding standards before any code is accepted into the repository.
- D. Require developers to run automated integration tests on their local development environments before release.
- E. Configure a CI server. Add a suite of unit tests to your code and have your CI server run them on commit and verify any changes.

Answer: B,E

Question: 60

[Building and implementing CI/CD pipelines for a service]

You are building an application that runs on Cloud Run. The application needs to access a third-party API by using an API key. You need to determine a secure way to store and use the API key in your application by following Google-recommended practices. What should you do?

- A. Save the API key in Secret Manager as a secret. Reference the secret as an environment variable in the Cloud Run application.
- B. Save the API key in Secret Manager as a secret key. Mount the secret key under the `/sys/api_key` directory and decrypt the key in the Cloud Run application.
- C. Save the API key in Cloud Key Management Service (Cloud KMS) as a key. Reference the key as an environment variable in the Cloud Run application.
- D. Encrypt the API key by using Cloud Key Management Service (Cloud KMS) and pass the key to Cloud Run as an environment variable. Decrypt and use the key in Cloud Run.

Answer: A

Question: 61

[Building and implementing CI/CD pipelines for a service]

You are implementing a CI/CD pipeline for your application in your company's multi-cloud environment. Your application is deployed by using custom Compute Engine images and the equivalent in other cloud providers. You need to implement a solution that will enable you to build and deploy the images to your current environment and is adaptable to future changes. Which solution stack should you use?

- A. Cloud Build with Packer
- B. Cloud Build with Google Cloud Deploy
- C. Google Kubernetes Engine with Google Cloud Deploy
- D. Cloud Build with kpt

Answer: B

Question: 62

[Building and implementing CI/CD pipelines for a service]

Your organization has a containerized web application that runs on-premises. As part of the migration plan to Google Cloud, you need to select a deployment strategy and platform that meets the following acceptance criteria:

1. The platform must be able to direct traffic from Android devices to an Android-specific microservice.
 2. The platform must allow for arbitrary percentage-based traffic splitting.
 3. The deployment strategy must allow for continuous testing of multiple versions of any microservice.
- What should you do?

- A. Deploy the canary release of the application to Cloud Run. Use traffic splitting to direct 10% of user traffic to the canary release based on the revision tag.
- B. Deploy the canary release of the application to App Engine. Use traffic splitting to direct a subset of user traffic to the new version based on the IP address.
- C. Deploy the canary release of the application to Compute Engine. Use Anthos Service Mesh with Compute Engine to direct 10% of user traffic to the canary release by configuring the virtual service.
- D. Deploy the canary release to Google Kubernetes Engine with Anthos Service Mesh. Use traffic splitting to direct 10% of user traffic to the new version based on the user-agent header configured in the virtual service.

Answer: D

Question: 63

[Applying Site Reliability Engineering (SRE) principles to a service]

Your company experiences bugs, outages, and slowness in its production systems. Developers use the production environment for new feature development and bug fixes. Configuration and experiments are done in the production environment, causing outages for users. Testers use the production environment for load testing, which often slows the production systems. You need to redesign the environment to reduce the number of bugs and outages in production and to enable testers to load test new features. What should you do?

- A. Create an automated testing script in production to detect failures as soon as they occur.
- B. Create a development environment with smaller server capacity and give access only to developers and testers.
- C. Secure the production environment to ensure that developers can't change it and set up one controlled update per year.
- D. Create a development environment for writing code and a test environment for configurations, experiments, and load testing.

Answer: D

Question: 64

[Building and implementing CI/CD pipelines for a service]

You have a CI/CD pipeline that uses Cloud Build to build new Docker images and push them to Docker Hub. You use Git for code versioning. After making a change in the Cloud Build YAML configuration, you notice that no new artifacts are being built by the pipeline. You need to resolve the issue following Site Reliability Engineering practices. What should you do?

- A. Disable the CI pipeline and revert to manually building and pushing the artifacts.
- B. Change the CI pipeline to push the artifacts to Container Registry instead of Docker Hub.
- C. Upload the configuration YAML file to Cloud Storage and use Error Reporting to identify and fix the issue.
- D. Run a Git compare between the previous and current Cloud Build Configuration files to find and fix the bug.

Answer: D

Question: 65

[Building and implementing service monitoring strategies]

You use Google Cloud Managed Service for Prometheus with managed collection to gather metrics from your service running on Google Kubernetes Engine (GKE). After deploying the service, there is no metric data appearing in Cloud Monitoring, and you have not encountered any error messages. You need to troubleshoot this issue. What should you do?

- A. Determine if your service has exceeded its quota for writes to the Cloud Monitoring API.
- B. Check if the Grafana service is installed on your GKE cluster.
- C. Confirm that your service has the monitoring.servicesViewer IAM role.
- D. Verify that your PodMonitoring configuration references a valid port.

Answer: D

Question: 66

[Building and implementing service monitoring strategies]

Your company runs services on Google Cloud. Each team runs their applications in a dedicated project. New teams and projects are created regularly. Your security team requires that all logs are processed by a security information and event management (SIEM) system. The SIEM ingests logs by using Pub/Sub. You must ensure that all existing and future logs are scanned by the SIEM. What should you do?

- A. Create an organization-level aggregated sink with a SIEM Pub/Sub topic as the destination. Set an inclusion filter to include all logs.

- B. Create a project-level logging sink with a SIEM Pub/Sub topic as the destination. Set an inclusion filter to include all logs. Repeat for each project.
- C. Create an organization-level aggregated sink with a SIEM log bucket as the destination. Set an inclusion filter to include all logs.
- D. Create a folder-level aggregated sink with a SIEM Pub/Sub topic as the destination. Set an inclusion filter to include all logs. Repeat for each folder.

Answer: A

Question: 67

[Building and implementing service monitoring strategies]

You are managing an application that exposes an HTTP endpoint without using a load balancer. The latency of the HTTP responses is important for the user experience. You want to understand what HTTP latencies all of your users are experiencing. You use Stackdriver Monitoring. What should you do?

- A. • In your application, create a metric with a metricKind set to DELTA and a valueType set to DOUBLE. • In Stackdriver's Metrics Explorer, use a Slacked Bar graph to visualize the metric.
- B. • In your application, create a metric with a metricKind set to CUMULATIVE and a valueType set to DOUBLE. • In Stackdriver's Metrics Explorer, use a Line graph to visualize the metric.
- C. • In your application, create a metric with a metricKind set to gauge and a valueType set to distribution. • In Stackdriver's Metrics Explorer, use a Heatmap graph to visualize the metric.
- D. • In your application, create a metric with a metricKind set to METRIC_KIND_UNSPECIFIED and a valueType set to INT64. • In Stackdriver's Metrics Explorer, use a Stacked Area graph to visualize the metric.

Answer: C

Question: 68

[Building and implementing CI/CD pipelines for a service]

Your application images are built and pushed to Google Container Registry (GCR). You want to build an automated pipeline that deploys the application when the image is updated while minimizing the development effort. What should you do?

- A. Use Cloud Build to trigger a Spinnaker pipeline.
- B. Use Cloud Pub/Sub to trigger a Spinnaker pipeline.
- C. Use a custom builder in Cloud Build to trigger a Jenkins pipeline.
- D. Use Cloud Pub/Sub to trigger a custom deployment service running in Google Kubernetes Engine (GKE).

Answer: B

Question: 69

[Building and implementing CI/CD pipelines for a service]

You support an application deployed on Compute Engine. The application connects to a Cloud SQL instance to store and retrieve data.

- a. After an update to the application, users report errors showing database timeout messages. The number of concurrent active users remained stable. You need to find the most probable cause of the database timeout. What should you do?
- A. Check the serial port logs of the Compute Engine instance.
 - B. Use Stackdriver Profiler to visualize the resources utilization throughout the application.
 - C. Determine whether there is an increased number of connections to the Cloud SQL instance.
 - D. Use Cloud Security Scanner to see whether your Cloud SQL is under a Distributed Denial of Service (DDoS) attack.

Answer: C

Question: 70

[Optimizing service performance]

You manage several production systems that run on Compute Engine in the same Google Cloud Platform (GCP) project. Each system has its own set of dedicated Compute Engine instances. You want to know how much it costs to run each of the systems. What should you do?

- A. In the Google Cloud Platform Console, use the Cost Breakdown section to visualize the costs per system.
- B. Assign all instances a label specific to the system they run. Configure BigQuery billing export and query costs per label.
- C. Enrich all instances with metadata specific to the system they run. Configure Stackdriver Logging to export to BigQuery, and query costs based on the metadata.
- D. Name each virtual machine (VM) after the system it runs. Set up a usage report export to a Cloud Storage bucket. Configure the bucket as a source in BigQuery to query costs based on VM name.

Answer: B

Question: 71

[Building and implementing service monitoring strategies]

You are running a real-time gaming application on Compute Engine that has a production and testing environment. Each environment has their own Virtual Private Cloud (VPC) network. The application frontend and backend servers are located on different subnets in the environment's VPC. You suspect there is a malicious process communicating intermittently in your production frontend servers. You want to ensure that network traffic is captured for analysis. What should you do?

- A. Enable VPC Flow Logs on the production VPC network frontend and backend subnets only with a sample volume scale of 0.5.
- B. Enable VPC Flow Logs on the production VPC network frontend and backend subnets only with a sample volume scale of 1.0.
- C. Enable VPC Flow Logs on the testing and production VPC network frontend and backend subnets with a volume scale of 0.5. Apply changes in testing before production.
- D. Enable VPC Flow Logs on the testing and production VPC network frontend and backend subnets with a volume scale of 1.0. Apply changes in testing before production.

Answer: D

Question: 72

[Building and implementing CI/CD pipelines for a service]

You are performing a semiannual capacity planning exercise for your flagship service. You expect a service user growth rate of 10% month-over-month over the next six months. Your service is fully containerized and runs on Google Cloud Platform (GCP), using a Google Kubernetes Engine (GKE) Standard regional cluster on three zones with cluster autoscaler enabled. You currently consume about 30% of your total deployed CPU capacity, and you require resilience against the failure of a zone. You want to ensure that your users experience minimal negative impact as a result of this growth or as a result of zone failure, while avoiding unnecessary costs. How should you prepare to handle the predicted growth?

- A. Verify the maximum node pool size, enable a horizontal pod autoscaler, and then perform a load test to verify your expected resource needs.
- B. Because you are deployed on GKE and are using a cluster autoscaler, your GKE cluster will scale automatically, regardless of growth rate.
- C. Because you are at only 30% utilization, you have significant headroom and you won't need to add any additional capacity for this rate of growth.
- D. Proactively add 60% more node capacity to account for six months of 10% growth rate, and then perform a load test to make sure you have enough capacity.

Answer: A

Question: 73

[Building and implementing service monitoring strategies]

You manage an application that is writing logs to Stackdriver Logging. You need to give some team members the ability to export logs. What should you do?

- A. Grant the team members the IAM role of logging.configWriter on Cloud IAM.
- B. Configure Access Context Manager to allow only these members to export logs.
- C. Create and grant a custom IAM role with the permissions logging.sinks.list and logging.sink.get.
- D. Create an Organizational Policy in Cloud IAM to allow only these members to create log exports.

Answer: A

Question: 74

[Building and implementing CI/CD pipelines for a service]

You built a serverless application by using Cloud Run and deployed the application to your production environment. You want to identify the resource utilization of the application for cost optimization. What should you do?

- A. Use Cloud Trace with distributed tracing to monitor the resource utilization of the application
- B. Use Cloud Profiler with Ops Agent to monitor the CPU and memory utilization of the application
- C. Use Cloud Monitoring to monitor the container CPU and memory utilization of the application
- D. Use Cloud Ops to create logs-based metrics to monitor the resource utilization of the application

Answer: A

Question: 75

[Applying Site Reliability Engineering (SRE) principles to a service]

You are the Operations Lead for an ongoing incident with one of your services. The service usually runs at around 70% capacity. You notice that one node is returning 5xx errors for all requests. There has also been a noticeable increase in support cases from customers. You need to remove the offending node from the load balancer pool so that you can isolate and investigate the node. You want to follow Google-recommended practices to manage the incident and reduce the impact on users. What should you do?

- A. 1. Communicate your intent to the incident team. 2. Perform a load analysis to determine if the remaining nodes can handle the increase in traffic offloaded from the removed node, and scale appropriately. 3. When any new nodes report healthy, drain traffic from the unhealthy node, and remove the unhealthy node from service.
- B. 1. Communicate your intent to the incident team. 2. Add a new node to the pool, and wait for the new node to report as healthy. 3. When traffic is being served on the new node, drain traffic from the unhealthy node, and remove the old node from service.
- C. 1. Drain traffic from the unhealthy node and remove the node from service. 2. Monitor traffic to ensure that the error is resolved and that the other nodes in the pool are handling the traffic appropriately. 3. Scale the pool as necessary to handle the new load. 4. Communicate your actions to the incident team.
- D. 1. Drain traffic from the unhealthy node and remove the old node from service. 2. Add a new node to the pool, wait for the new node to report as healthy, and then serve traffic to the new node. 3. Monitor traffic to ensure that the pool is healthy and is handling traffic appropriately. 4. Communicate your actions to the incident team.

Answer: A

Question: 76

[Optimizing service performance]

You are working with a government agency that requires you to archive application logs for seven years. You need to configure Stackdriver to export and store the logs while minimizing costs of storage. What should you do?

- A. Create a Cloud Storage bucket and develop your application to send logs directly to the bucket.
- B. Develop an App Engine application that pulls the logs from Stackdriver and saves them in BigQuery.
- C. Create an export in Stackdriver and configure Cloud Pub/Sub to store logs in permanent storage for seven years.
- D. Create a sink in Stackdriver, name it, create a bucket on Cloud Storage for storing archived logs, and then select the bucket as the log export destination.

Answer: D

Question: 77

[Building and implementing service monitoring strategies]

Your company has recently experienced several production service issues. You need to create a Cloud Monitoring dashboard to troubleshoot the issues, and you want to use the dashboard to distinguish between failures in your own service and those caused by a Google Cloud service that you use. What should you do?

- A. Enable Personalized Service Health annotations on the dashboard.
- B. Create an alerting policy for the system error metrics.
- C. Create a log-based metric to track cloud service errors, and display the metric on the dashboard.
- D. Create a logs widget to display system errors from Cloud Logging on the dashboard.

Answer: A

Question: 78

[Building and implementing service monitoring strategies]

You are configuring Cloud Logging for a new application that runs on a Compute Engine instance with a public IP address. A user-managed service account is attached to the instance. You confirmed that the necessary agents are running on the instance but you cannot see any log entries from the instance in Cloud Logging. You want to resolve the issue by following Google-recommended practices. What should you do?

- A. Add the Logs Writer role to the service account.
- B. Enable Private Google Access on the subnet that the instance is in.
- C. Update the instance to use the default Compute Engine service account.
- D. Export the service account key and configure the agents to use the key.

Answer: A

Question: 79

[Building and implementing service monitoring strategies]

A third-party application needs to have a service account key to work properly. When you try to export the key from your cloud project, you receive an error: "The organization policy constraint `iam.disableServiceAccountKeyCreation` is enforced. You need to make the third-party application work while following Google-recommended security practices. What should you do?

- A. Enable the default service account key, and download the key.
- B. Remove the `iam.disableServiceAccountKeyCreation` policy at the organization level, and create a key.
- C. Disable the service account key creation policy at the project's folder, and download the default key.
- D. Add a rule to set the `iam.disableServiceAccountKeyCreation` policy to off in your project and create a key.

Answer: D

Question: 80

[Building and implementing CI/CD pipelines for a service]

You manage an application that runs in Google Kubernetes Engine (GKE) and uses the blue/green deployment methodology. Extracts of the Kubernetes manifests are shown below:

```
apiVersion: apps/v1 kind: Deployment metadata:
```

```
  name: app-green labels:
```

```
    app: my-app version: green
```

```
<?other fields snippet>
```

```
apiVersion: v1 kind: Service metadata:
```

```
  name: app-svc spec:
```

```
    selector:
```

```
      app: my-app version: green
```

```
<?other fields snippet>
```

```
apiVersion: networking.k8s.io/v1 kind: Ingress
```

```
metadata:
```

The Deployment app-green was updated to use the new version of the application. During post-deployment monitoring, you notice that the majority of user requests are failing. You did not observe this behavior in the testing environment. You need to mitigate the incident impact on users and enable the developers to troubleshoot the issue.

What should you do?

- A. Update the Deployment app-blue to use the new version of the application.
- B. Update the Deployment app-green to use the previous version of the application.
- C. Change the selector on the Service app-svc to app: my-app.
- D. Change the selector on the Service app-svc to app: my-app, version: blue.

Answer: C

Question: 81

[Applying Site Reliability Engineering (SRE) principles to a service]

You have an application running in Google Kubernetes Engine. The application invokes multiple services per request but responds too slowly. You need to identify which downstream service or services are causing the delay. What should you do?

- A. Analyze VPC flow logs along the path of the request.
- B. Investigate the Liveness and Readiness probes for each service.
- C. Create a Dataflow pipeline to analyze service metrics in real time.
- D. Use a distributed tracing framework such as OpenTelemetry or Stackdriver Trace.

Answer: C

Question: 82

[Building and implementing CI/CD pipelines for a service]

Your company operates in a highly regulated domain. Your security team requires that only trusted container images can be deployed to Google Kubernetes Engine (GKE). You need to implement a solution that meets the requirements of the security team, while minimizing management overhead. What should you do?

- A. Grant the roles/artifactregistry.writer role to the Cloud Build service account. Confirm that no employee has Artifact Registry write permission.
- B. Use Cloud Run to write and deploy a custom validator. Enable an Eventarc trigger to perform validations when new images are uploaded.
- C. Configure Kritis to run in your GKE clusters to enforce deploy-time security policies.
- D. Configure Binary Authorization in your GKE clusters to enforce deploy-time security policies.

Answer: D

Question: 83

[Applying Site Reliability Engineering (SRE) principles to a service]

Your team is running microservices in Google Kubernetes Engine (GKE). You want to detect consumption of an error budget to protect customers and define release policies. What should you do?

- A. Create SLIs from metrics. Enable Alert Policies if the services do not pass.
- B. Use the metrics from Anthos Service Mesh to measure the health of the microservices.
- C. Create a SLO. Create an Alert Policy on `select_slo_burn_rate`.
- D. Create a SLO and configure uptime checks for your services. Enable Alert Policies if the services do not pass.

Answer: C

Question: 84

[Optimizing service performance]

You support the backend of a mobile phone game that runs on a Google Kubernetes Engine (GKE) cluster. The application is serving HTTP requests from users. You need to implement a solution that will reduce the network cost. What should you do?

- A. Configure the VPC as a Shared VPC Host project.
- B. Configure your network services on the Standard Tier.
- C. Configure your Kubernetes cluster as a Private Cluster.
- D. Configure a Google Cloud HTTP Load Balancer as Ingress.

Answer: B

Question: 85

[Building and implementing CI/CD pipelines for a service]

Your organization uses a change advisory board (CAB) to approve all changes to an existing service. You want to revise this process to eliminate any negative impact on the software delivery performance. What should you do?

Choose 2 answers

- A. Replace the CAB with a senior manager to ensure continuous oversight from development to deployment
- B. Let developers merge their own changes but ensure that the team's deployment platform can roll back changes if any issues are discovered
- C. Move to a peer-review based process for individual changes that is enforced at code check-in time and supported by automated tests
- D. Batch changes into larger but less frequent software releases
- E. Ensure that the team's development platform enables developers to get fast feedback on the impact of their changes

Answer: C,E

Question: 86

[Applying Site Reliability Engineering (SRE) principles to a service]

You need to introduce postmortems into your organization. You want to ensure that the postmortem process is well received. What should you do?

Choose 2 answers

- A. Create a designated team that is responsible for conducting all postmortems.
- B. Encourage new employees to conduct postmortems to learn through practice.
- C. Ensure that writing effective postmortems is a rewarded and celebrated practice.
- D. Encourage your senior leadership to acknowledge and participate in postmortems.
- E. Provide your organization with a forum to critique previous postmortems.

Answer: C,D

Question: 87

[Applying Site Reliability Engineering (SRE) principles to a service]

Your company follows Site Reliability Engineering principles. You are writing a postmortem for an incident, triggered by a software change, that severely affected users. You want to prevent severe incidents from happening in the future. What should you do?

- A. Identify engineers responsible for the incident and escalate to their senior management.
- B. Ensure that test cases that catch errors of this type are run successfully before new software releases.
- C. Follow up with the employees who reviewed the changes and prescribe practices they should follow in the future.
- D. Design a policy that will require on-call teams to immediately call engineers and management to discuss a plan of action if an incident occurs.

Answer: B

Question: 88

[Building and implementing CI/CD pipelines for a service]

You are deploying an application to Cloud Run. The application requires a password to start. Your organization requires that all passwords are rotated every 24 hours, and your application must have the latest password. You need to deploy the application with no downtime. What should you do?

- A. Store the password in Secret Manager and send the secret to the application by using environment variables.
- B. Store the password in Secret Manager and mount the secret as a volume within the application.
- C. Use Cloud Build to add your password into the application container at build time. Ensure that Artifact Registry is secured from public access.
- D. Store the password directly in the code. Use Cloud Build to rebuild and deploy the application each time the password changes.

Answer: B

Question: 89

[Building and implementing CI/CD pipelines for a service]

You use Cloud Build to build your application. You want to reduce the build time while minimizing cost and development effort. What should you do?

- A. Use Cloud Storage to cache intermediate artifacts.
- B. Run multiple Jenkins agents to parallelize the build.
- C. Use multiple smaller build steps to minimize execution time.
- D. Use larger Cloud Build virtual machines (VMs) by using the machine-type option.

Answer: C

Question: 90

[Building and implementing CI/CD pipelines for a service]

Your organization recently adopted a container-based workflow for application development. Your team develops numerous applications that are deployed continuously through an automated build pipeline to the production environment. A recent security audit alerted your team that the code pushed to production could contain vulnerabilities and that the existing tooling around virtual machine (VM) vulnerabilities no longer applies to the containerized environment. You need to ensure the security and patch level of all code running through the pipeline. What should you do?

- A. Set up Container Analysis to scan and report Common Vulnerabilities and Exposures.
- B. Configure the containers in the build pipeline to always update themselves before release.
- C. Reconfigure the existing operating system vulnerability software to exist inside the container.
- D. Implement static code analysis tooling against the Docker files used to create the containers.

Answer: D

Question: 91

[Building and implementing CI/CD pipelines for a service]

Your company runs applications in Google Kubernetes Engine (GKE) that are deployed following a GitOps methodology.

Application developers frequently create cloud resources to support their applications. You want to give developers the ability to manage infrastructure as code, while ensuring that you follow Google- recommended practices. You need to ensure that infrastructure as code reconciles periodically to avoid configuration drift. What should you do?

- A. Install and configure Config Connector in Google Kubernetes Engine (GKE).
- B. Configure Cloud Build with a Terraform builder to execute plan and apply commands.
- C. Create a Pod resource with a Terraform docker image to execute terraform plan and terraform apply commands.
- D. Create a Job resource with a Terraform docker image to execute terraform plan and terraform apply commands.

Answer: A

Question: 92

[Building and implementing CI/CD pipelines for a service]

You are using Jenkins to orchestrate your CI/CD pipelines deploying your applications to GKE and onpremises VMs. After expanding your on-premises infrastructure, all resource utilization is normal, but you notice that initiating on-premises deployments is taking longer than expected. You need to accelerate the initiation of on-premises deployments. What should you do?

- A. Implement a blue-green deployment strategy for your on-premises VMs.
- B. Upgrade your on-premises VMs to have faster CPUs and more memory.
- C. Increase the number of Jenkins executor threads.
- D. Configure Jenkins to use an artifact repository in your on-premises environment.

Answer: D

Question: 93

[Building and implementing service monitoring strategies]

As a Site Reliability Engineer, you support an application written in GO that runs on Google Kubernetes Engine (GKE) in production. After releasing a new version Of the application, you notice the application runs for about 15 minutes and then restarts. You decide to add Cloud Profiler to your application and now notice that the heap usage grows constantly until the application restarts. What should you do?

- A. Add high memory compute nodes to the cluster.
- B. Increase the memory limit in the application deployment.
- C. Add Cloud Trace to the application, and redeploy.
- D. Increase the CPU limit in the application deployment.

Answer: B

Question: 94

[Building and implementing CI/CD pipelines for a service]

Your team uses Cloud Build for all CI/CO pipelines. You want to use the kubectl builder for Cloud Build to deploy new images to Google Kubernetes Engine (GKE). You need to authenticate to GKE while minimizing development effort. What should you do?

- A. Assign the Container Developer role to the Cloud Build service account.
- B. Specify the Container Developer role for Cloud Build in the cloudbuild.yaml file.
- C. Create a new service account with the Container Developer role and use it to run Cloud Build.
- D. Create a separate step in Cloud Build to retrieve service account credentials and pass these to kubectl.

Answer: A

Question: 95

[Building and implementing service monitoring strategies]

You are using Stackdriver to monitor applications hosted on Google Cloud Platform (GCP). You recently deployed a new application, but its logs are not appearing on the Stackdriver dashboard.

You need to troubleshoot the issue. What should you do?

- A. Confirm that the Stackdriver agent has been installed in the hosting virtual machine.
- B. Confirm that your account has the proper permissions to use the Stackdriver dashboard.
- C. Confirm that port 25 has been opened in the firewall to allow messages through to Stackdriver.
- D. Confirm that the application is using the required client library and the service account key has proper permissions.

Answer: A

Question: 96

[Building and implementing service monitoring strategies]

You have deployed a fleet of Compute Engine instances in Google Cloud. You need to ensure that monitoring metrics and logs for the instances are visible in Cloud Logging and Cloud Monitoring by your company's operations and cyber

security teams. You need to grant the required roles for the Compute Engine service account by using Identity and Access Management (IAM) while following the principle of least privilege. What should you do?

- A. Grant the logging.editor and monitoring.metricwriter roles to the Compute Engine service accounts.
- B. Grant the Logging, admin and monitoring . editor roles to the Compute Engine service accounts.
- C. Grant the logging, logwriter and monitoring, editor roles to the Compute Engine service accounts.
- D. Grant the logging, logWriter and monitoring, metricWriter roles to the Compute Engine service accounts.

Answer: A

Question: 97

[Managing service incidents]

Your company follows Site Reliability Engineering practices. You are the person in charge of Communications for a large, ongoing incident affecting your customer-facing applications. There is still no estimated time for a resolution of the outage. You are receiving emails from internal stakeholders who want updates on the outage, as well as emails from customers who want to know what is happening. You want to efficiently provide updates to everyone affected by the outage. What should you do?

- A. Focus on responding to internal stakeholders at least every 30 minutes. Commit to "next update" times.
- B. Provide periodic updates to all stakeholders in a timely manner. Commit to a "next update" time in all communications.
- C. Delegate the responding to internal stakeholder emails to another member of the Incident Response Team. Focus on providing responses directly to customers.
- D. Provide all internal stakeholder emails to the Incident Commander, and allow them to manage internal communications. Focus on providing responses directly to customers.

Answer: B

Question: 98

[Building and implementing service monitoring strategies]

You support a trading application written in Python and hosted on App Engine flexible environment.

You want to customize the error information being sent to Stackdriver Error Reporting. What should you do?

- A. Install the Stackdriver Error Reporting library for Python, and then run your code on a Compute Engine VM.
- B. Install the Stackdriver Error Reporting library for Python, and then run your code on Google Kubernetes Engine.
- C. Install the Stackdriver Error Reporting library for Python, and then run your code on App Engine flexible environment.
- D. Use the Stackdriver Error Reporting API to write errors from your application to `ReportedErrorEvent`, and then generate log entries with properly formatted error messages in Stackdriver Logging.

Answer: D

Question: 99

[Building and implementing service monitoring strategies]

Your team of Infrastructure DevOps Engineers is growing, and you are starting to use Terraform to manage infrastructure. You need a way to implement code versioning and to share code with other team members. What should you do?

- A. Store the Terraform code in a version-control system. Establish procedures for pushing new versions and merging with the master.
- B. Store the Terraform code in a network shared folder with child folders for each version release. Ensure that everyone works on different files.
- C. Store the Terraform code in a Cloud Storage bucket using object versioning. Give access to the bucket to every team member so they can download the files.
- D. Store the Terraform code in a shared Google Drive folder so it syncs automatically to every team member's computer. Organize files with a naming convention that identifies each new version.

Answer: A

Question: 100

[Building and implementing CI/CD pipelines for a service]

You are performing a semi-annual capacity planning exercise for your flagship service. You expect a service user growth rate of 10% month-over-month for the next six months. Your service is fully containerized and runs on a Google Kubernetes Engine (GKE) standard cluster across three zones with cluster autoscaling enabled. You currently consume about 30% of your total deployed CPU capacity and you require resilience against the failure of a zone. You want to ensure that your users experience minimal negative impact as a result of this growth or as a result of zone failure while you avoid unnecessary costs. How should you prepare to handle the predicted growth?

- A. Verify the maximum node pool size, enable a Horizontal Pod Autoscaler, and then perform a load test to verify your expected resource needs.
- B. Because you deployed the service on GKE and are using a cluster autoscaler, your GKE cluster will scale automatically regardless of growth rate.
- C. Because you are only using 30% of deployed CPU capacity, there is significant headroom and you do not need to add any additional capacity for this rate of growth.
- D. Proactively add 80% more node capacity to account for six months of 10% growth rate and then perform a load test to ensure that you have enough capacity.

Answer: A

Question: 101

[Applying Site Reliability Engineering (SRE) principles to a service]

Your organization wants to implement Site Reliability Engineering (SRE) culture and principles. Recently, a service that you support had a limited outage. A manager on another team asks you to provide a formal explanation of what happened so they can action remediations. What should you do?

- A. Develop a postmortem that includes the root causes, resolution, lessons learned, and a prioritized list of action items. Share it with the manager only.
- B. Develop a postmortem that includes the root causes, resolution, lessons learned, and a prioritized list of action items. Share it on the engineering organization's document portal.

- C. Develop a postmortem that includes the root causes, resolution, lessons learned, the list of people responsible, and a list of action items for each person. Share it with the manager only.
- D. Develop a postmortem that includes the root causes, resolution, lessons learned, the list of people responsible, and a list of action items for each person. Share it on the engineering organization's document portal.

Answer: B

Question: 102

[Building and implementing CI/CD pipelines for a service]

You are designing a system with three different environments: development, quality assurance (QA), and production. Each environment will be deployed with Terraform and has a Google Kubernetes Engine Enterprise (GKE Enterprise) cluster created so that application teams can deploy their applications. Config Sync will be used and templated to deploy infrastructure-level resources in each GKE Enterprise cluster. All users (for example, infrastructure operators and application owners) will use GitOps. How should you structure your source control repositories for both infrastructure as code (IaC) and application code?

- A. Cloud Infrastructure (Terraform) repository is shared: different directories are different environmentsGKE Enterprise Infrastructure (Config Sync Kustomize manifests) repository is shared: different overlay directories are different environmentsApplication (app source code) repositories are separated: different branches are different features
- B. Cloud Infrastructure (Terraform) repository is shared: different directories are different environmentsGKE Enterprise Infrastructure (Config Sync Kustomize manifests) repositories are separated: different branches are different environmentsApplication (app source code) repositories are separated: different branches are different features
- C. Cloud Infrastructure (Terraform) repository is shared: different branches are different environmentsGKE Enterprise Infrastructure (Config Sync Kustomize manifests) repository is shared: different overlay directories are different environmentsApplication (app source code) repository is shared: different directories are different features
- D. Cloud Infrastructure (Terraform) repositories are separated: different branches are different environmentsGKE Enterprise Infrastructure (Config Sync Kustomize manifests) repositories are separated: different overlay directories are different environmentsApplication (app source code) repositories are separated: different branches are different features

Answer: A

Question: 103

[Building and implementing CI/CD pipelines for a service]

You have an application that runs on Cloud Run. You want to use live production traffic to test a new version of the application while you let the quality assurance team perform manual testing. You want to limit the potential impact of any issues while testing the new version, and you must be able to roll back to a previous version of the application if needed. How should you deploy the new version? Choose 2 answers

- A. Deploy the application as a new Cloud Run service.
- B. Deploy a new Cloud Run revision with a tag and use the —no-traffic option.

- C. Deploy a new Cloud Run revision without a tag and use the `—no-traffic` option.
- D. Deploy the new application version and use the `—no-traffic` option Route production traffic to the revision's URL.
- E. Deploy the new application version and split traffic to the new version.

Answer: B,E

Question: 104

[Building and implementing service monitoring strategies]

You are configuring connectivity across Google Kubernetes Engine (GKE) clusters in different VPCs You notice that the nodes in Cluster A are unable to access the nodes in Cluster B You suspect that the workload access issue is due to the network configuration You need to troubleshoot the issue but do not have execute access to workloads and nodes You want to identify the layer at which the network connectivity is broken What should you do?

- A. Install a toolbox container on the node in Cluster A Confirm that the routes to Cluster B are configured appropriately
- B. Use Network Connectivity Center to perform a Connectivity Test from Cluster A to Cluster
- C. Use a debug container to run the `tracert` command from Cluster A to Cluster B and from Cluster B to Cluster A Identify the common failure point
- D. Enable VPC Flow Logs in both VPCs and monitor packet drops

Answer: C

Question: 105

[Building and implementing CI/CD pipelines for a service]

You are troubleshooting a failed deployment in your CI/CD pipeline. The deployment logs indicate that the application container failed to start due to a missing environment variable. You need to identify the root cause and implement a solution within your CI/CD workflow to prevent this issue from recurring. What should you do?

- A. Run integration tests in the CI pipeline.
- B. Implement static code analysis in the CI pipeline.
- C. Use a canary deployment strategy.
- D. Enable Cloud Audit Logs for the deployment.

Answer: A

Question: 106

[Building and implementing CI/CD pipelines for a service]

You are designing a system with three different environments: development, quality assurance (QA), and production.

Each environment will be deployed with Terraform and has a Google Kubernetes Engine (GKE) cluster created so that application teams can deploy their applications. Anthos Config Management will be used and templated to deploy

infrastructure level resources in each GKE cluster. All users (for example, infrastructure operators and application owners) will use GitOps. How should you structure your source control repositories for both Infrastructure as Code (IaC) and application code?

- A. Cloud Infrastructure (Terraform) repository is shared: different directories are different environmentsGKE Infrastructure (Anthos Config Management Kustomize manifests) repository is shared: different overlay directories are different environmentsApplication (app source code) repositories are separated: different branches are different features
- B. Cloud Infrastructure (Terraform) repository is shared: different directories are different environmentsGKE Infrastructure (Anthos Config Management Kustomize manifests) repositories are separated: different branches are different environmentsApplication (app source code) repositories are separated: different branches are different features
- C. Cloud Infrastructure (Terraform) repository is shared: different branches are different environmentsGKE Infrastructure (Anthos Config Management Kustomize manifests) repository is shared: different overlay directories are different environmentsApplication (app source code) repository is shared: different directories are different features
- D. Cloud Infrastructure (Terraform) repositories are separated: different branches are different environmentsGKE Infrastructure (Anthos Config Management Kustomize manifests) repositories are separated: different overlay directories are different environmentsApplication (app source code) repositories are separated: different branches are different features

Answer: B

Question: 107

[Applying Site Reliability Engineering (SRE) principles to a service]

You are leading a DevOps project for your organization. The DevOps team is responsible for managing the service infrastructure and being on-call for incidents. The Software Development team is responsible for writing, submitting, and reviewing code. Neither team has any published SLOs. You want to design a new joint-ownership model for a service between the DevOps team and the Software Development team. Which responsibilities should be assigned to each team in the new joint-ownership model?

DevOps team responsibilities	Software Development team responsibilities
• Manage the service infrastructure • Be on-call for incidents • Perform code reviews	• Submit code to be reviewed by the DevOps team • Publish the SLOs that the DevOps team must meet

DevOps team responsibilities	Software Development team responsibilities
• Manage the service infrastructure • Perform code reviews	• Submit code to be reviewed by the DevOps team • Be on-call for incidents • Publish the SLOs that the DevOps team must meet

DevOps team responsibilities	Shared responsibilities	Software Development team responsibilities
• Manage the service infrastructure	• Perform code reviews • Be on-call for incidents on a rotation basis • Adopt and publish SLOs for the service	• Submit code to be reviewed

D.

DevOps team responsibilities	Shared responsibilities	Software Development team responsibilities
• Manage the service infrastructure • Be on-call for incidents	• Adopt and publish SLOs for the service	• Submit code to be reviewed • Perform code reviews

A. Option A B. Option B C. Option C D. Option D

Answer: D

Question: 108

[Building and implementing CI/CD pipelines for a service]

You are building and running client applications in Cloud Run and Cloud Functions Your client requires that all logs must be available for one year so that the client can import the logs into their logging service You must minimize required code changes What should you do?

A. Update all images in Cloud Run and all functions in Cloud Functions to send logs to both Cloud Logging and the client's logging service Ensure that all the ports required to send logs are open in the VPC firewall

- B. Create a Pub/Sub topic subscription and logging sink. Configure the logging sink to send all logs into the topic. Give your client access to the topic to retrieve the logs.
- C. Create a storage bucket and appropriate VPC firewall rules. Update all images in Cloud Run and all functions in Cloud Functions to send logs to a file within the storage bucket.
- D. Create a logs bucket and logging sink. Set the retention on the logs bucket to 365 days. Configure the logging sink to send logs to the bucket. Give your client access to the bucket to retrieve the logs.

Answer: D

Question: 109

[Building and implementing CI/CD pipelines for a service]

Your application runs on Google Cloud Platform (GCP). You need to implement Jenkins for deploying application releases to GCP. You want to streamline the release process, lower operational toil, and keep user data secure. What should you do?

- A. Implement Jenkins on local workstations.
- B. Implement Jenkins on Kubernetes on-premises.
- C. Implement Jenkins on Google Cloud Functions.
- D. Implement Jenkins on Compute Engine virtual machines.

Answer: D

Question: 110

[Building and implementing CI/CD pipelines for a service]

Your company is using HTTPS requests to trigger a public Cloud Run-hosted service accessible at the `https://booking-engine-abcdef.a.run.app` URL. You need to give developers the ability to test the latest revisions of the service before the service is exposed to customers. What should you do?

- A. `Run the gcloud run deploy booking-engine --no-traffic --ag dev` command. Use the `https://dev-booking-engine-abcdef.a.run.app` URL for testing.
- B. `Run the gcloud run services update-traffic booking-engine --to-revisions LATEST*!` command. Use the `https://booking-engine-abcdef.a.run.app` URL for testing.
- C. Pass the `curl -K "Authorization: Bearer $(gcloud auth print-identity-token)"` auth token. Use the `https://booking-engine-abcdef.a.run.app` URL to test privately.
- D. Grant the roles/run.invoker role to the developers testing the booking-engine service. Use the `https://booking-engine-abcdef.private.run.app` URL for testing.

Answer: B

Question: 111

[Optimizing service performance]

You support a Node.js application running on Google Kubernetes Engine (GKE) in production. The application makes several HTTP requests to dependent applications. You want to anticipate which dependent applications might cause performance issues. What should you do?

- A. Instrument all applications with Stackdriver Profiler.
- B. Instrument all applications with Stackdriver Trace and review inter-service HTTP requests.
- C. Use Stackdriver Debugger to review the execution of logic within each application to instrument all applications.
- D. Modify the Node.js application to log HTTP request and response times to dependent applications. Use Stackdriver Logging to find dependent applications that are performing poorly.

Answer: B

Question: 112

[Building and implementing service monitoring strategies]

You are creating and assigning action items in a postmortem for an outage. The outage is over, but you need to address the root causes. You want to ensure that your team handles the action items quickly and efficiently. How should you assign owners and collaborators to action items?

- A. Assign one owner for each action item and any necessary collaborators.
- B. Assign multiple owners for each item to guarantee that the team addresses items quickly
- C. Assign collaborators but no individual owners to the items to keep the postmortem blameless.
- D. Assign the team lead as the owner for all action items because they are in charge of the SRE team.

Answer: A

Question: 113

[Building and implementing service monitoring strategies]

You are running an application on Compute Engine and collecting logs through Stackdriver. You discover that some personally identifiable information (PII) is leaking into certain log entry fields. You want to prevent these fields from being written in new log entries as quickly as possible. What should you do?

- A. Use the filter-record-transformer Fluentd filter plugin to remove the fields from the log entries in flight.
- B. Use the fluent-plugin-record-reformer Fluentd output plugin to remove the fields from the log entries in flight.
- C. Wait for the application developers to patch the application, and then verify that the log entries are no longer

exposing PII.

D. Stage log entries to Cloud Storage, and then trigger a Cloud Function to remove the fields and write the entries to Stackdriver via the Stackdriver Logging API.

Answer: A

Question: 114

[Building and implementing service monitoring strategies]

You support an application running on GCP and want to configure SMS notifications to your team for the most critical alerts in Stackdriver Monitoring. You have already identified the alerting policies you want to configure this for.

What should you do?

- A. Download and configure a third-party integration between Stackdriver Monitoring and an SMS gateway. Ensure that your team members add their SMS/phone numbers to the external tool.
- B. Select the Webhook notifications option for each alerting policy, and configure it to use a third-party integration tool. Ensure that your team members add their SMS/phone numbers to the external tool.
- C. Ensure that your team members set their SMS/phone numbers in their Stackdriver Profile. Select the SMS notification option for each alerting policy and then select the appropriate SMS/phone numbers from the list.
- D. Configure a Slack notification for each alerting policy. Set up a Slack-to-SMS integration to send SMS messages when Slack messages are received. Ensure that your team members add their SMS/phone numbers to the external integration.

Answer: C

Question: 115

[Building and implementing CI/CD pipelines for a service]

You are ready to deploy a new feature of a web-based application to production. You want to use Google Kubernetes Engine (GKE) to perform a phased rollout to half of the web server pods.

What should you do?

- A. Use a partitioned rolling update.
- B. Use Node taints with NoExecute.
- C. Use a replica set in the deployment specification.
- D. Use a stateful set with parallel pod management policy.

Answer: A

Question: 116

[Building and implementing CI/CD pipelines for a service]

You are running a web application that connects to an AlloyDB cluster by using a private IP address in your default VPC. You need to run a database schema migration in your CI/CD pipeline by using Cloud Build before deploying a new version of your application. You want to follow Google- recommended security practices. What should you do?

- A. Set up a Cloud Build private pool to access the database through a static external IP address. Configure the database to only allow connections from this IP address. Execute the schema migration script in the private pool.
- B. Create a service account that has permission to access the database. Configure Cloud Build to use this service account and execute the schema migration script in a private pool.
- C. Add the database username and encrypted password to the application configuration file. Use these credentials in Cloud Build to execute the schema migration script.
- D. Add the database username and password to Secret Manager. When running the schema migration script, retrieve the username and password from Secret Manager.

Answer: B

Question: 117

[Building and implementing service monitoring strategies]

You have a set of applications running on a Google Kubernetes Engine (GKE) cluster, and you are using Stackdriver Kubernetes Engine Monitoring. You are bringing a new containerized application required by your company into production. This application is written by a third party and cannot be modified or reconfigured. The application writes its log information to `/var/log/app_messages.log`, and you want to send these log entries to Stackdriver Logging. What should you do?

- A. Use the default Stackdriver Kubernetes Engine Monitoring agent configuration.
- B. Deploy a Fluentd daemonset to GKE. Then create a customized input and output configuration to tail the log file in the application's pods and write to Stackdriver Logging.
- C. Install Kubernetes on Google Compute Engine (GCE) and redeploy your applications. Then customize the built-in Stackdriver Logging configuration to tail the log file in the application's pods and write to Stackdriver Logging.
- D. Write a script to tail the log file within the pod and write entries to standard output. Run the script as a sidecar container with the application's pod. Configure a shared volume between the containers to allow the script to have read access to `/var/log` in the application container.

Answer: B

Question: 118

[Building and implementing CI/CD pipelines for a service]

Some of your production services are running in Google Kubernetes Engine (GKE) in the eu-west-1 region. Your build system runs in the us-west-1 region. You want to push the container images from your build system to a scalable registry to maximize the bandwidth for transferring the images to the cluster. What should you do?

- A. Push the images to Google Container Registry (GCR) using the gcr.io hostname.
- B. Push the images to Google Container Registry (GCR) using the us.gcr.io hostname.
- C. Push the images to Google Container Registry (GCR) using the eu.gcr.io hostname.
- D. Push the images to a private image registry running on a Compute Engine instance in the eu-west-1 region.

Answer: C

Question: 119

[Building and implementing CI/CD pipelines for a service]

You are managing the production deployment to a set of Google Kubernetes Engine (GKE) clusters.

You want to make sure only images which are successfully built by your trusted CI/CD pipeline are deployed to production. What should you do?

- A. Enable Cloud Security Scanner on the clusters.
- B. Enable Vulnerability Analysis on the Container Registry.
- C. Set up the Kubernetes Engine clusters as private clusters.
- D. Set up the Kubernetes Engine clusters with Binary Authorization.

Answer: D

Question: 120

[Applying Site Reliability Engineering (SRE) principles to a service]

You recently noticed that one Of your services has exceeded the error budget for the current rolling window period.

Your company's product team is about to launch a new feature. You want to follow Site Reliability Engineering (SRE) practices.

What should you do?

- A. Notify the team that their error budget is used up. Negotiate with the team for a launch freeze or tolerate a

slightly worse user experience.

B. Look through other metrics related to the product and find SLOs with remaining error budget.

Reallocate the error budgets and allow the feature launch.

C. Escalate the situation and request additional error budget.

D. Notify the team about the lack of error budget and ensure that all their tests are successful so the launch will not further risk the error budget.

Answer: A

Question: 121

[Building and implementing service monitoring strategies]

You are running an application in a virtual machine (VM) using a custom Debian image. The image has the Stackdriver Logging agent installed. The VM has the cloud-platform scope. The application is logging information via syslog. You want to use Stackdriver Logging in the Google Cloud Platform Console to visualize the logs. You notice that syslog is not showing up in the "All logs" dropdown list of the Logs Viewer. What is the first thing you should do?

A. Look for the agent's test log entry in the Logs Viewer.

B. Install the most recent version of the Stackdriver agent.

C. Verify the VM service account access scope includes the monitoring.write scope.

D. SSH to the VM and execute the following commands on your VM: `ps ax | grep fluentd`

Answer: D

Question: 122

[Optimizing service performance]

You need to reduce the cost of virtual machines (VM) for your organization. After reviewing different options, you decide to leverage preemptible VM instances. Which application is suitable for preemptible VMs?

A. A scalable in-memory caching system

B. The organization's public-facing website

C. A distributed, eventually consistent NoSQL database cluster with sufficient quorum

D. A GPU-accelerated video rendering platform that retrieves and stores videos in a storage bucket

Answer: D

Question: 123

[Building and implementing service monitoring strategies]

Your company recently migrated to Google Cloud. You need to design a fast, reliable, and repeatable solution for your company to provision new projects and basic resources in Google Cloud. What should you do?

- A. Use the Google Cloud console to create projects.
- B. Write a script by using the gcloud CLI that passes the appropriate parameters from the request. Save the script in a Git repository.
- C. Write a Terraform module and save it in your source control repository. Copy and run the apply command to create the new project.
- D. Use the Terraform repositories from the Cloud Foundation Toolkit. Apply the code with appropriate parameters to create the Google Cloud project and related resources.

Answer: D

Question: 124

[Building and implementing service monitoring strategies]

You recently configured an App Hub application. You are able to see the managed instance group, backend service, and URL map listed in App Hub, but you do not see the forwarding rule. You must ensure that the forwarding rule is listed. What should you do?

- A. Attach the project containing the forwarding rule as an App Hub service project.
- B. Enable the App Hub API in the project containing the forwarding rule.
- C. Configure the forwarding rule to forward to the correct target proxy.
- D. Register the forwarding rule as a service in the application configuration.

Answer: D

Question: 125

[Building and implementing service monitoring strategies]

You are managing an application that runs in Compute Engine. The application uses a custom HTTP server to expose an API that is accessed by other applications through an internal TCP/UDP load balancer. A firewall rule allows access to the API port from 0.0.0.0/0. You need to configure Cloud Logging to log each IP address that accesses the API by using the fewest number of steps. What should you do, Bret?

- A. Enable Packet Mirroring on the VPC.
- B. Install the Ops Agent on the Compute Engine instances.
- C. Enable logging on the firewall rule.
- D. Enable VPC Flow Logs on the subnet.

Answer: D

Question: 126

[Building and implementing CI/CD pipelines for a service]

Your team deploys applications to three Google Kubernetes Engine (GKE) environments development staging and production You use GitHub repositories as your source of truth You need to ensure that the three environments are consistent You want to follow Google-recommended practices to enforce and install network policies and a logging DaemonSet on all the GKE clusters in those environments What should you do?

- A. Use Google Cloud Deploy to deploy the network policies and the DaemonSet Use Cloud Monitoring to trigger an alert if the network policies and DaemonSet drift from your source in the repository.
- B. Use Google Cloud Deploy to deploy the DaemonSet and use Policy Controller to configure the network policies Use Cloud Monitoring to detect drifts from the source in the repository and Cloud Functions to correct the drifts
- C. Use Cloud Build to render and deploy the network policies and the DaemonSet Set up Config Sync to sync the configurations for the three environments
- D. Use Cloud Build to render and deploy the network policies and the DaemonSet Set up a Policy Controller to enforce the configurations for the three environments

Answer: C

Question: 127

[Building and implementing CI/CD pipelines for a service]

Your product is currently deployed in three Google Cloud Platform (GCP) zones with your users divided between the zones. You can fail over from one zone to another, but it causes a 10-minute service disruption for the affected users. You typically experience a database failure once per quarter and can detect it within five minutes. You are cataloging the reliability risks of a new real-time chat feature for your product. You catalog the following information for each risk:

- Mean Time to Detect (MTTD) in minutes
- Mean Time to Repair (MTTR) in minutes
- Mean Time Between Failure (MTBF) in days
- User Impact Percentage

The chat feature requires a new database system that takes twice as long to successfully fail over between zones. You want to account for the risk of the new database failing in one zone. What would be the values for the risk of database failover with the new system?

- A. MTTD: 5MTTR: 10MTBF: 90Impact: 33%
- B. MTTD: 5MTTR: 20MTBF: 90Impact: 33%
- C. MTTD: 5MTTR: 10MTBF: 90Impact: 50%
- D. MTTD: 5MTTR: 20MTBF: 90Impact: 50%

Answer: B

Question: 128

[Building and implementing CI/CD pipelines for a service]

Your team has recently deployed an NGINX-based application into Google Kubernetes Engine (GKE) and has exposed it to the public via an HTTP Google Cloud Load Balancer (GCLB) ingress. You want to scale the deployment of the

application's frontend using an appropriate Service Level Indicator (SLI). What should you do?

- A. Configure the horizontal pod autoscaler to use the average response time from the Liveness and Readiness probes.
- B. Configure the vertical pod autoscaler in GKE and enable the cluster autoscaler to scale the cluster as pods expand.
- C. Install the Stackdriver custom metrics adapter and configure a horizontal pod autoscaler to use the number of requests provided by the GCLB.
- D. Expose the NGINX stats endpoint and configure the horizontal pod autoscaler to use the request metrics exposed by the NGINX deployment.

Answer: C

Question: 129

[Applying Site Reliability Engineering (SRE) principles to a service]

You support a large service with a well-defined Service Level Objective (SLO). The development team deploys new releases of the service multiple times a week. If a major incident causes the service to miss its SLO, you want the development team to shift its focus from working on features to improving service reliability. What should you do before a major incident occurs?

- A. Develop an appropriate error budget policy in cooperation with all service stakeholders.
- B. Negotiate with the product team to always prioritize service reliability over releasing new features.
- C. Negotiate with the development team to reduce the release frequency to no more than once a week.
- D. Add a plugin to your Jenkins pipeline that prevents new releases whenever your service is out of SLO.

Answer: A

Question: 130

[Building and implementing service monitoring strategies]

You need to create a Cloud Monitoring SLO for a service that will be published soon. You want to verify that requests to the service will be addressed in fewer than 300 ms at least 90% Of the time per calendar month. You need to identify the metric and evaluation method to use. What should you do?

- A. Select a latency metric for a request-based method of evaluation.
- B. Select a latency metric for a window-based method of evaluation.
- C. Select an availability metric for a request-based method of evaluation.
- D. Select an availability metric for a window-based method Of evaluation.

Answer: A

Question: 131

[Building and implementing service monitoring strategies]

You recently migrated an ecommerce application to Google Cloud. You now need to prepare the application for the upcoming peak traffic season. You want to follow Google-recommended practices. What should you do first to prepare for the busy season?

- A. Migrate the application to Cloud Run, and use autoscaling.
- B. Load test the application to profile its performance for scaling.
- C. Create a Terraform configuration for the application's underlying infrastructure to quickly deploy to additional regions.
- D. Pre-provision the additional compute power that was used last season, and expect growth.

Answer: B

Question: 132

[Applying Site Reliability Engineering (SRE) principles to a service]

You are leading a DevOps project for your organization. The DevOps team is responsible for managing the service infrastructure and being on-call for incidents. The Software Development team

is responsible for writing, submitting, and reviewing code. Neither team has any published SLOs. You want to design a new joint-ownership model for a service between the DevOps team and the Software Development team. Which responsibilities should be assigned to each team in the new joint-ownership model?

- A. DevOps team responsibilitiesManage the service infrastructureBe on-call for incidentsPerform code reviewsSoftware Development team responsibilitiesSubmit code to be reviewed by the DevOps teamPublish the SLOs that the DevOps team must meet
- B. DevOps team responsibilitiesManage the service infrastructurePerform code reviewsSoftware Development team responsibilitiesSubmit code to be reviewed by the DevOps teamBe on-call for incidentsPublish the SLOs that the DevOps team must meet
- C. DevOps team responsibilitiesShared responsibilities for code reviewsSoftware Development team responsibilitiesManage the service infrastructureBe on-call for incidents on a rotation basisAdopt and publish SLOs for the serviceSubmit code to be reviewed
- D. DevOps team responsibilitiesManage the service infrastructureBe on-call for incidentsSoftware Development team responsibilitiesAdopt and publish SLOs for the serviceSubmit code to be reviewedShared responsibilities for code reviews

Answer: D

Question: 133

[Applying Site Reliability Engineering (SRE) principles to a service]

You support a user-facing web application. When analyzing the application's error budget over the previous six months, you notice that the application never consumed more than 5% of its error budget. You hold a SLO review with business stakeholders and confirm that the SLO is set appropriately. You want your application's reliability to more closely reflect its SLO. What steps can you take to further that goal while balancing velocity, reliability, and business needs? Choose 2 answers.

- A. Add more serving capacity to all of your application's zones
- B. Implement and measure all other available SLIs for the application
- C. Announce planned downtime to consume more error budget and ensure that users are not depending on a tighter SLO
- D. Have more frequent or potentially risky application releases
- E. Tighten the SLO to match the application's observed reliability

Answer: D,E

Question: 134

[Applying Site Reliability Engineering (SRE) principles to a service]

You are running an experiment to see whether your users like a new feature of a web application.

Shortly after deploying the feature as a canary release, you receive a spike in the number of 500 errors sent to users, and your monitoring reports show increased latency. You want to quickly minimize the negative impact on users. What should you do first?

- A. Roll back the experimental canary release.
- B. Start monitoring latency, traffic, errors, and saturation.
- C. Record data for the postmortem document of the incident.
- D. Trace the origin of 500 errors and the root cause of increased latency.

Answer: A

Question: 135

[Building and implementing CI/CD pipelines for a service]

You work for a company that manages highly sensitive user data.

a. You are designing the Google Kubernetes Engine (GKE) infrastructure for your company, including several applications that will be deployed in development and production environments. Your design must protect data from unauthorized access from other applications while minimizing the amount of management overhead required.

What should you do?

- A. Create one cluster for the organization with separate namespaces for each application and environment combination.
- B. Create one cluster for each environment (development and production) with each application in its own namespace within each cluster.
- C. Create one cluster for the organization with separate namespaces for each application.
- D. Create one cluster for each application with separate namespaces for production and development environments.

Answer: B

Question: 136

[Applying Site Reliability Engineering (SRE) principles to a service]

You support a stateless web-based API that is deployed on a single Compute Engine instance in the europe-west2-a zone. The Service Level Indicator (SLI) for service availability is below the specified Service Level Objective (SLO). A postmortem has revealed that requests to the API regularly time out. The time outs are due to the API having a high number of requests and running out memory. You want to improve service availability. What should you do?

- A. Change the specified SLO to match the measured SLI.
- B. Move the service to higher-specification compute instances with more memory.
- C. Set up additional service instances in other zones and load balance the traffic between all instances.
- D. Set up additional service instances in other zones and use them as a failover in case the primary instance is unavailable.

Answer: C

Question: 137

[Applying Site Reliability Engineering (SRE) principles to a service]

You support a service with a well-defined Service Level Objective (SLO). Over the previous 6 months, your service has consistently met its SLO and customer satisfaction has been consistently high. Most of your service's operations tasks are automated and few repetitive tasks occur frequently. You want to optimize the balance between reliability and deployment velocity while following site reliability engineering best practices. What should you do? (Choose two.)

- A. Make the service's SLO more strict.
- B. Increase the service's deployment velocity and/or risk.
- C. Shift engineering time to other services that need more reliability.
- D. Get the product team to prioritize reliability work over new features.
- E. Change the implementation of your Service Level Indicators (SLIs) to increase coverage.

Answer: B,C

Question: 138

[Applying Site Reliability Engineering (SRE) principles to a service]

You support a high-traffic web application with a microservice architecture. The home page of the application displays multiple widgets containing content such as the current weather, stock prices, and news headlines. The main serving thread makes a call to a dedicated microservice for each widget and then lays out the homepage for the user. The microservices occasionally fail; when that happens, the serving thread serves the homepage with some missing content. Users of the application are unhappy if this degraded mode occurs too frequently, but they would rather have some content served instead of no content at all. You want to set a Service Level Objective (SLO) to ensure that the user experience does not degrade too much. What Service Level Indicator (SLI) should you use to measure this?

- A. A quality SLI: the ratio of non-degraded responses to total responses
- B. An availability SLI: the ratio of healthy microservices to the total number of microservices
- C. A freshness SLI: the proportion of widgets that have been updated within the last 10 minutes
- D. A latency SLI: the ratio of microservice calls that complete in under 100 ms to the total number of microservice calls

Answer: B

Question: 139

[Optimizing service performance]

Your application's performance in Google Cloud has degraded since the last release. You suspect that downstream dependencies might be causing some requests to take longer to complete. You need to investigate the issue with your application to determine the cause. What should you do?

- A. Configure Cloud Trace in your application.
- B. Configure Error Reporting in your application.
- C. Configure Cloud Profiler in your application.
- D. Configure Google Cloud Managed Service for Prometheus in your application.

Answer: A

Question: 140

[Building and implementing CI/CD pipelines for a service]

You are deploying a Cloud Build job that deploys Terraform code when a Git branch is updated. While testing, you noticed that the job fails. You see the following error in the build logs:

Initializing the backend. ..

Error: Failed to get existing workspaces : querying Cloud Storage failed: googleapi : Error

You need to resolve the issue by following Google-recommended practices. What should you do?

- A. Change the Terraform code to use local state.
- B. Create a storage bucket with the name specified in the Terraform configuration.
- C. Grant the roles/ owner Identity and Access Management (IAM) role to the Cloud Build service account on the project.
- D. Grant the roles/ storage. objectAdmin Identity and Access Management (IAM) role to the Cloud Build service account on the state file bucket.

Answer: D

Question: 141

[Building and implementing CI/CD pipelines for a service]

Your organization is using Helm to package containerized applications Your applications reference

both public and private charts Your security team flagged that using a public Helm repository as a dependency is a risk You want to manage all charts uniformly, with native access control and VPC Service Controls What should you do?

- A. Store public and private charts in OCI format by using Artifact Registry
- B. Store public and private charts by using GitHub Enterprise with Google Workspace as the identity provider
- C. Store public and private charts by using Git repository Configure Cloud Build to synchronize contents of the repository into a Cloud Storage bucket Connect Helm to the bucket by using `https://[bucket].storage.googleapis.com/[helmchart]` as the Helm repository
- D. Configure a Helm chart repository server to run in Google Kubernetes Engine (GKE) with Cloud Storage bucket as the storage backend

Answer: A

Question: 142

[Building and implementing CI/CD pipelines for a service]

Your organization recently adopted a container-based workflow for application development. Your team develops numerous applications that are deployed continuously through an automated build pipeline to a Kubernetes cluster in the production environment. The security auditor is concerned that developers or operators could circumvent automated testing and push code changes to production without approval. What should you do to enforce approvals?

- A. Configure the build system with protected branches that require pull request approval.
- B. Use an Admission Controller to verify that incoming requests originate from approved sources.
- C. Leverage Kubernetes Role-Based Access Control (RBAC) to restrict access to only approved users.
- D. Enable binary

authorization inside the Kubernetes cluster and configure the build pipeline as an attestor.

Answer: D

Question: 143

[Optimizing service performance]

You need to run a business-critical workload on a fixed set of Compute Engine instances for several months. The workload is stable with the exact amount of resources allocated to it. You want to lower the costs for this workload without any performance implications. What should you do?

- A. Purchase Committed Use Discounts.
- B. Migrate the instances to a Managed Instance Group.
- C. Convert the instances to preemptible virtual machines.
- D. Create an Unmanaged Instance Group for the instances used to run the workload.

Answer: A

Question: 144

[Building and implementing CI/CD pipelines for a service]

You use Terraform to manage an application deployed to a Google Cloud environment. The application runs on instances deployed by a managed instance group. The Terraform code is deployed by using a CI/CD pipeline. When you change the machine type on the instance template used by the managed instance group, the pipeline fails at the terraform apply stage with the following error message:

Error waiting for Deleting Instance Template: The instance_template resource 'projects/my-project/global/instanceTemplates/my-it-2C22C101C1CT01C1i . 001' is already being used by 'projects/my-project/regions/ns-aentrall/instanceGroupManagars/niy- mig'

You need to update the instance template and minimize disruption to the application and the number of pipeline runs. What should you do?

- A. Delete the managed instance group and recreate it after updating the instance template.
- B. Add a new instance template, update the managed instance group to use the new instance template, and delete the old instance template.
- C. Remove the managed instance group from the Terraform state file, update the instance template, and reimport the managed instance group.
- D. Set the `create_before_destroy` meta-argument to true in the lifecycle block on the instance template.

Answer: D

Question: 145

[Building and implementing CI/CD pipelines for a service]

Your company is developing applications that are deployed on Google Kubernetes Engine (GKE). Each team manages a different application. You need to create the development and production environments for each team while you minimize costs. Different teams should not be able to access other teams' environments. You want to follow Google-recommended practices. What should you do?

- A. Create one Google Cloud project per team. In each project, create a cluster for development and one for production. Grant the teams Identity and Access Management (IAM) access to their respective clusters.
- B. Create one Google Cloud project per team. In each project, create a cluster with a Kubernetes namespace for development and one for production. Grant the teams Identity and Access Management (IAM) access to their respective clusters.
- C. Create a development and a production GKE cluster in separate projects. In each cluster, create a Kubernetes namespace per team and then configure Identity-Aware Proxy so that each team can only access its own namespace.
- D. Create a development and a production GKE cluster in separate projects. In each cluster, create a Kubernetes namespace per team and then configure Kubernetes role-based access control (RBAC) so that each team can only access its own namespace.

Answer: D

Question: 146

[Applying Site Reliability Engineering (SRE) principles to a service]

You support a popular mobile game application deployed on Google Kubernetes Engine (GKE) across several Google Cloud regions. Each region has multiple Kubernetes clusters. You receive a report that none of the users in a specific region can connect to the application. You want to resolve the incident while following Site Reliability Engineering practices. What should you do first?

- A. Reroute the user traffic from the affected region to other regions that don't report issues.
- B. Use Stackdriver Monitoring to check for a spike in CPU or memory usage for the affected region.
- C. Add an extra node pool that consists of high memory and high CPU machine type instances to the cluster.
- D. Use Stackdriver Logging to filter on the clusters in the affected region, and inspect error messages in the logs.

Answer: A

Question: 147

[Building and implementing CI/CD pipelines for a service]

You are developing the deployment and testing strategies for your CI/CD pipeline in Google Cloud. You must be able to

- E. Reduce the complexity of release deployments and minimize the duration of deployment rollbacks
- F. Test real production traffic with a gradual increase in the number of affected users

You want to select a deployment and testing strategy that meets your requirements. What should you do?

- A. Recreate deployment and canary testing
- B. Blue/green deployment and canary testing
- C. Rolling update deployment and A/B testing
- D. Rolling update deployment and shadow testing

Answer: B

Question: 148

[Managing service incidents]

You are on-call for an infrastructure service that has a large number of dependent systems. You receive an alert indicating that the service is failing to serve most of its requests and all of its dependent systems with hundreds of thousands of users are affected. As part of your Site Reliability Engineering (SRE) incident management protocol, you declare yourself Incident Commander (IC) and pull in two experienced people from your team as Operations Lead (OL) and Communications Lead (CL). What should you do next?

- A. Look for ways to mitigate user impact and deploy the mitigations to production.
- B. Contact the affected service owners and update them on the status of the incident.
- C. Establish a communication channel where incident responders and leads can communicate with each other.
- D. Start a postmortem, add incident information, circulate the draft internally, and ask internal stakeholders for input.

Answer: A

Question: 149

[Building and implementing CI/CD pipelines for a service]

Your team is building a service that performs compute-heavy processing on batches of data. The data is processed faster based on the speed and number of CPUs on the machine. These batches of data vary in size and may arrive at any time from multiple third-party sources. You need to ensure that third parties are able to upload their data securely.

You want to minimize costs while ensuring that the data is processed as quickly as possible. What should you do?

- A. • Provide a secure file transfer protocol (SFTP) server on a Compute Engine instance so that third parties can upload batches of data and provide appropriate credentials to the server • Create a Cloud Function with a google.storage, object, finalize Cloud Storage trigger. Write code so that the function can scale up a Compute Engine autoscaling managed instance group • Use an image preloaded with the data processing software that terminates the instances when processing completes
- B. • Provide a Cloud Storage bucket so that third parties can upload batches of data, and provide appropriate Identity and Access Management (IAM) access to the bucket • Use a standard

Google Kubernetes Engine (GKE) cluster and maintain two services one that processes the batches of data and one that monitors Cloud Storage for new batches of data • Stop the processing service when there are no batches of data to process

C. • Provide a Cloud Storage bucket so that third parties can upload batches of data, and provide appropriate identity and Access Management (IAM) access to the bucket • Create a Cloud Function with a google, storage, object .finalise Cloud Storage trigger Write code so that the function can scale up a Compute Engine autoscaling managed instance group • Use an image pre-loaded with the data

processing software that terminates the instances when processing completes

D. • Provide a Cloud Storage bucket so that third parties can upload batches of data, and provide appropriate Identity and Access Management (IAM) access to the bucket • Use Cloud Monitoring to detect new batches of data in the bucket and trigger a Cloud Function that processes the data • Set a Cloud Function to use the largest CPU possible to minimize the runtime of the processing

Answer: C

Question: 150

[Building and implementing service monitoring strategies]

Your company's security team needs to have read-only access to Data Access audit logs in the _Required bucket You want to provide your security team with the necessary permissions following the principle of least privilege and Google-recommended practices. What should you do?

- A. Assign the roles/logging.viewer role to each member of the security team
- B. Assign the roles/logging.viewer role to a group with all the security team members
- C. Assign the roles/logging.privateLogViewer role to each member of the security team
- D. Assign the roles/logging.privateLogviewer role to a group with all the security team members

Answer: D

Question: 151

[Building and implementing service monitoring strategies]

Your company runs services by using Google Kubernetes Engine (GKE). The GKE clusters in the development environment run applications with verbose logging enabled. Developers view logs by using the kubectl logs command and do not use Cloud Logging. Applications do not have a uniform logging structure defined. You need to minimize the costs associated with application logging while still collecting GKE operational logs. What should you do?

- A. Run the gcloud container clusters update --logging=SYSTEM command for the development cluster.
- B. Run the gcloud container clusters update logging=WORKLOAD command for the development cluster.
- C. Run the gcloud logging sinks update _Default --disabled command in the project associated with the development environment.

D. Add the severity $\geq$ DEBUG resource. type "k83 container" exclusion filter to the Default logging sink in the project associated with the development environment.

Answer: A

Question: 152

[Optimizing service performance]

You are monitoring a service that uses n2-standard-2 Compute Engine instances that serve large files. Users have reported that downloads are slow. Your Cloud Monitoring dashboard shows that your VMS are running at peak network throughput. You want to improve the network throughput performance. What should you do?

- A. Deploy a Cloud NAT gateway and attach the gateway to the subnet of the VMS.
- B. Add additional network interface controllers (NICs) to your VMS.
- C. Change the machine type for your VMS to n2-standard-8.
- D. Deploy the Ops Agent to export additional monitoring metrics.

Answer: C

Question: 153

[Building and implementing CI/CD pipelines for a service]

Your company is developing applications that are deployed on Google Kubernetes Engine (GKE). Each team manages a different application. You need to create the development and production environments for each team, while minimizing costs. Different teams should not be able to access other teams' environments. What should you do?

- A. Create one GCP Project per team. In each project, create a cluster for Development and one for Production. Grant the teams IAM access to their respective clusters.
- B. Create one GCP Project per team. In each project, create a cluster with a Kubernetes namespace for Development and one for Production. Grant the teams IAM access to their respective clusters.
- C. Create a Development and a Production GKE cluster in separate projects. In each cluster, create a Kubernetes namespace per team, and then configure Identity Aware Proxy so that each team can only access its own namespace.
- D. Create a Development and a Production GKE cluster in separate projects. In each cluster, create a Kubernetes namespace per team, and then configure Kubernetes Role-based access control (RBAC) so that each team can only access its own namespace.

Answer: D

Question: 154

[Building and implementing service monitoring strategies]

Your company allows teams to self-manage Google Cloud projects, including project-level Identity and Access Management (IAM). You are concerned that the team responsible for the Shared VPC project might accidentally delete the project, so a lien has been placed on the project. You need to design a solution to restrict Shared VPC project deletion to those with the `resourceManager.projects.updateLiens` permission at the organization level. What should you do?

- A. Enable VPC Service Controls for the container.googleapis.com API service.
- B. Revoke the `resourceManager.projects.updateLiens` permission from all users associated with the project.
- C. Enable the `compute.restrictXpnProjectLienRemoval` organization policy constraint.
- D. Instruct teams to only perform IAM permission management as code with Terraform.

Answer: C

Question: 155

[Building and implementing service monitoring strategies]

Your company runs applications in Google Kubernetes Engine (GKE). Several applications rely on ephemeral volumes. You noticed some applications were unstable due to the `DiskPressure` node condition on the worker nodes. You need to identify which Pods are causing the issue, but you do not have execute access to workloads and nodes. What should you do?

- A. Check the `node/ephemeral_storage/used_bytes` metric by using Metrics Explorer.
- B. Check the metric by using Metrics Explorer.
- C. Locate all the Pods with `emptyDir` volumes. Use the `df-h` command to measure volume disk usage.
- D. Locate all the Pods with `emptyDir` volumes. Use the `du -sh *` command to measure volume disk usage.

Answer: A

Question: 156

[Building and implementing CI/CD pipelines for a service]

You are configuring your CI/CD pipeline natively on Google Cloud. You want builds in a preproduction Google Kubernetes Engine (GKE) environment to be automatically load-tested before being promoted to the production GKE environment. You need to ensure that only builds that have passed this test are deployed to production. You want to follow Google-recommended practices.

How should you configure this pipeline with Binary Authorization?

- A. Create an attestation for the builds that pass the load test by requiring the lead quality assurance engineer to sign the attestation by using a key stored in Cloud Key Management Service (Cloud KMS).

- B. Create an attestation for the builds that pass the load test by using a private key stored in Cloud Key Management Service (Cloud KMS) authenticated through Workload Identity.
- C. Create an attestation for the builds that pass the load test by using a private key stored in Cloud Key Management Service (Cloud KMS) with a service account JSON key stored as a Kubernetes Secret.
- D. Create an attestation for the builds that pass the load test by requiring the lead quality assurance engineer to sign the attestation by using their personal private key.

Answer: B

Question: 157

[Applying Site Reliability Engineering (SRE) principles to a service]

You support a high-traffic web application and want to ensure that the home page loads in a timely manner. As a first step, you decide to implement a Service Level Indicator (SLI) to represent home page request latency with an acceptable page load time set to 100 ms. What is the Google- recommended way of calculating this SLI?

- A. Bucketize the request latencies into ranges, and then compute the percentile at 100 ms.
- B. Bucketize the request latencies into ranges, and then compute the median and 90th percentiles.
- C. Count the number of home page requests that load in under 100 ms, and then divide by the total number of home page requests.
- D. Count the number of home page requests that load in under 100 ms. and then divide by the total number of all web application requests.

Answer: C

Question: 158

[Building and implementing service monitoring strategies]

You are developing a strategy for monitoring your Google Cloud Platform (GCP) projects in production using Stackdriver Workspaces. One of the requirements is to be able to quickly identify and react to production environment issues without false alerts from development and staging projects. You want to ensure that you adhere to the principle of least privilege when providing relevant team members with access to Stackdriver Workspaces. What should you do?

- A. Grant relevant team members read access to all GCP production projects. Create Stackdriver workspaces inside each project.
- B. Grant relevant team members the Project Viewer IAM role on all GCP production projects. Create Stackdriver workspaces inside each project.
- C. Choose an existing GCP production project to host the monitoring workspace. Attach the production projects to this workspace. Grant relevant team members read access to the Stackdriver Workspace.
- D. Create a new GCP monitoring project, and create a Stackdriver Workspace inside it. Attach the production projects to this workspace. Grant relevant team members read access to the Stackdriver Workspace.

Answer: D

Question: 159

[Building and implementing CI/CD pipelines for a service]

Your application services run in Google Kubernetes Engine (GKE). You want to make sure that only images from your centrally-managed Google Container Registry (GCR) image registry in the altostrat- images project can be deployed to the cluster while minimizing development time. What should you do?

- A. Create a custom builder for Cloud Build that will only push images to gcr.io/altostrat-images.
- B. Use a Binary Authorization policy that includes the whitelist name pattern gcr.io/altostrat- images/.
- C. Add logic to the deployment pipeline to check that all manifests contain only images from gcr.io/altostrat-images.
- D. Add a tag to each image in gcr.io/altostrat-images and check that this tag is present when the image is deployed.

Answer: B

Question: 160

[Optimizing service performance]

You work for a global organization and are running a monolithic application on Compute Engine. You need to select the machine type for the application to use that optimizes CPU utilization by using the fewest number of steps. You want to use historical system metrics to identify the machine type for the application to use. You want to follow Google-recommended practices. What should you do?

- A. Use the Recommender API and apply the suggested recommendations.
- B. Create an Agent Policy to automatically install Ops Agent in all VMs.
- C. Install the Ops Agent in a fleet of VMs by using the gcloud CLI.
- D. Review the Cloud Monitoring dashboard for the VM and choose the machine type with the lowest CPU utilization.

Answer: A

Question: 161

[Optimizing service performance]

Your company runs an ecommerce website built with JVM-based applications and microservice architecture in Google Kubernetes Engine (GKE). The application load increases during the day and decreases during the night. Your operations team has configured the application to run enough Pods to handle the evening peak load. You want to automate scaling by only running enough Pods and nodes for the load. What should you do?

- A. Configure the Vertical Pod Autoscaler but keep the node pool size static.
- B. Configure the Vertical Pod Autoscaler and enable the cluster autoscaler.

- C. Configure the Horizontal Pod Autoscaler but keep the node pool size static
- D. Configure the Horizontal Pod Autoscaler and enable the cluster autoscaler

Answer: D

Question: 162

[Building and implementing CI/CD pipelines for a service]

You use Cloud Build to build and deploy your application. You want to securely incorporate database credentials and other application secrets into the build pipeline. You also want to minimize the development effort. What should you do?

- A. Create a Cloud Storage bucket and use the built-in encryption at rest. Store the secrets in the bucket and grant Cloud Build access to the bucket.
- B. Encrypt the secrets and store them in the application repository. Store a decryption key in a separate repository and grant Cloud Build access to the repository.
- C. Use client-side encryption to encrypt the secrets and store them in a Cloud Storage bucket. Store a decryption key in the bucket and grant Cloud Build access to the bucket.
- D. Use Cloud Key Management Service (Cloud KMS) to encrypt the secrets and include them in your Cloud Build deployment configuration. Grant Cloud Build access to the KeyRing.

Answer: D

Question: 163

[Building and implementing service monitoring strategies]

You support a high-traffic web application that runs on Google Cloud Platform (GCP). You need to

measure application reliability from a user perspective without making any engineering changes to it. What should you do?

Choose 2 answers

- A. Review current application metrics and add new ones as needed.
- B. Modify the code to capture additional information for user interaction.
- C. Analyze the web proxy logs only and capture response time of each request.
- D. Create new synthetic clients to simulate a user journey using the application.
- E. Use current and historic Request Logs to trace customer interaction with the application.

Answer: D,E

Question: 164

[Building and implementing CI/CD pipelines for a service]

Your development team has created a new version of their service's API. You need to deploy the new versions of the API with the least disruption to third-party developers and end users of third-party installed applications. What should you do?

- A. Introduce the new version of the API. Announce deprecation of the old version of the API. Deprecate the old version of the API. Contact remaining users of the old API. Provide best effort support to users of the old API. Turn down the old version of the API.
- B. Announce deprecation of the old version of the API. Introduce the new version of the API. Contact remaining users on the old API. Deprecate the old version of the API. Turn down the old version of the API. Provide best effort support to users of the old API.
- C. Announce deprecation of the old version of the API. Contact remaining users on the old API. Introduce the new version of the API. Deprecate the old version of the API. Provide best effort support to users of the old API. Turn down the old version of the API.
- D. Introduce the new version of the API. Contact remaining users of the old API. Announce deprecation of the old version of the API. Deprecate the old version of the API. Turn down the old version of the API. Provide best effort support to users of the old API.

Answer: A

Question: 165

[Managing service incidents]

Your company follows Site Reliability Engineering practices. You are the Incident Commander for a new, customer-impacting incident. You need to immediately assign two incident management roles to assist you in an effective incident response. What roles should you assign?

Choose 2 answers

- A. Operations Lead
- B. Engineering Lead
- C. Communications Lead
- D. Customer Impact Assessor
- E. External Customer Communications Lead

Answer: A,C

Question: 166

[Building and implementing service monitoring strategies]

Your team is preparing to launch a new API in Cloud Run. The API uses an OpenTelemetry agent to send distributed tracing data to Cloud Trace to monitor the time each request takes. The team has noticed inconsistent trace collection.

You need to resolve the issue. What should you do?

- A. Increase the CPU limit in Cloud Run from 2 to 4.
- B. Use an HTTP health check.
- C. Configure CPU to be allocated only during request processing.
- D. Configure CPU to be always-allocated.

Answer: D

Question: 167

[Building and implementing CI/CD pipelines for a service]

You need to build a CI/CD pipeline for a containerized application in Google Cloud Your development team uses a central Git repository for trunk-based development You want to run all your tests in the pipeline for any new versions of the application to improve the quality What should you do?

- A. 1. Install a Git hook to require developers to run unit tests before pushing the code to a central repository2. Trigger Cloud Build to build the application container Deploy the application container to a testing environment, and run integration tests3. If the integration tests are successful deploy the application container to your production environment. and run acceptance tests
- B. 1. Install a Git hook to require developers to run unit tests before pushing the code to a central repositoryIf all tests are successful build a container2. Trigger Cloud Build to deploy the application container to a testing environment, and run integrationtests and acceptance tests3. If all tests are successful tag the code as production ready Trigger Cloud Build to build and deploy the application container to the production environment
- C. 1. Trigger Cloud Build to build the application container and run unit tests with the container2. If unit tests are successful, deploy the application container to a testing environment, and run integration tests3. If the integration tests are successful the pipeline deploys the application container to the production environment After that, run acceptance tests
- D. 1. Trigger Cloud Build to run unit tests when the code is pushed If all unit tests are successful, build and push the application container to a central registry2. Trigger Cloud Build to deploy the container to a testing environment, and run integration tests and acceptance tests3. If all tests are successful the pipeline deploys the application to the production environment and runs smoke tests

Answer: D

Question: 168

[Building and implementing CI/CD pipelines for a service]

You use a multiple step Cloud Build pipeline to build and deploy your application to Google Kubernetes Engine (GKE). You want to integrate with a third-party monitoring platform by performing a HTTP POST of the build information to a webhook. You want to minimize the development effort. What should you do?

- A. Add logic to each Cloud Build step to HTTP POST the build information to a webhook.
- B. Add a new step at the end of the pipeline in Cloud Build to HTTP POST the build information to a webhook.
- C. Use Stackdriver Logging to create a logs-based metric from the Cloud Build logs. Create an Alert with a Webhook notification type.
- D. Create a Cloud Pub/Sub push subscription to the Cloud Build cloud-builds PubSub topic to HTTP POST the build information to a webhook.

Answer: D

Question: 169

[Applying Site Reliability Engineering (SRE) principles to a service]

You manage your company's primary revenue-generating application. You have an error budget policy in place that freezes production deployments when the application is close to breaching its SLO. A number of issues have recently occurred, and the application has exhausted its error budget. You need to deploy a new release to the application that includes a feature urgently required by your largest customer. You have been told that the release has passed all unit tests. What should you do?

- A. Start the deployment of the feature immediately.
- B. Delay the deployment of the feature until the error budget is replenished.
- C. Re-run the unit tests, and start the deployment of the feature if the tests pass.
- D. Deploy the feature to a subset of users, and gradually roll out to all users if there are no errors reported.

Answer: D

Question: 170

[Building and implementing service monitoring strategies]

Your team is designing a new application for deployment into Google Kubernetes Engine (GKE). You need to set up monitoring to collect and aggregate various application-level metrics in a centralized location. You want to use Google Cloud Platform services while minimizing the amount of work required to set up monitoring. What should you do?

- A. Publish various metrics from the application directly to the Stackdriver Monitoring API, and then observe these

custom metrics in Stackdriver.

- B. Install the Cloud Pub/Sub client libraries, push various metrics from the application to various topics, and then observe the aggregated metrics in Stackdriver.
- C. Install the OpenTelemetry client libraries in the application, configure Stackdriver as the export destination for the metrics, and then observe the application's metrics in Stackdriver.
- D. Emit all metrics in the form of application-specific log messages, pass these messages from the containers to the Stackdriver logging collector, and then observe metrics in Stackdriver.

Answer: A

Question: 171

[Building and implementing CI/CD pipelines for a service]

You are designing a deployment technique for your applications on Google Cloud. As part Of your deployment planning, you want to use live traffic to gather performance metrics for new versions Ofyour applications. You need to test against the full production load before your applications are launched. What should you do?

- A. Use A/B testing with blue/green deployment.
- B. Use shadow testing with continuous deployment.
- C. Use canary testing with continuous deployment.
- D. Use canary testing with rolling updates deployment,

Answer: B

Question: 172

[Building and implementing CI/CD pipelines for a service]

You need to enforce several constraint templates across your Google Kubernetes Engine (GKE) clusters. The constraints include policy parameters, such as restricting the Kubernetes API. You must ensure that the policy parameters are stored in a GitHub repository and automatically applied when changes occur. What should you do?

- A. Set up a GitHub action to trigger Cloud Build when there is a parameter change. In Cloud Build, run a gcloud CLI command to apply the change.
- B. When there is a change in GitHub, use a webhook to send a request to Cloud Service Mesh, and apply the change.
- C. Configure Config Sync with the GitHub repository. When there is a change in the repository, use Config Sync to apply the change.
- D. Configure Config Connector with the GitHub repository. When there is a change in the repository, use Config Connector to apply the change.

Answer: C

Question: 173

[Building and implementing CI/CD pipelines for a service]

You are using Terraform to manage infrastructure as code within a CI/CD pipeline. You notice that multiple copies of the entire infrastructure stack exist in your Google Cloud project, and a new copy is created each time a change to the existing infrastructure is made. You need to optimize your cloud spend by ensuring that only a single instance of your infrastructure stack exists at a time. You want to follow Google-recommended practices. What should you do?

- A. Create a new pipeline to delete old infrastructure stacks when they are no longer needed
- B. Confirm that the pipeline is storing and retrieving the terraform. tfstate file from Cloud Storage with the Terraform gcs backend
- C. Verify that the pipeline is storing and retrieving the terraform.tfstate* file from a source control
- D. Update the pipeline to remove any existing infrastructure before you apply the latest configuration

Answer: D

Question: 174

[Building and implementing CI/CD pipelines for a service]

You are creating a CI/CD pipeline to perform Terraform deployments of Google Cloud resources. Your CI/CD tooling is running in Google Kubernetes Engine (GKE) and uses an ephemeral Pod for each pipeline run. You must ensure that the pipelines that run in the Pods have the appropriate Identity and Access Management (IAM) permissions to perform the Terraform deployments. You want to follow Google-recommended practices for identity management. What should you do?

Choose 2 answers

- A. Create a new Kubernetes service account, and assign the service account to the Pods. Use Workload Identity to authenticate as the Google service account
- B. Create a new JSON service account key for the Google service account, store the key as a Kubernetes secret, inject the key into the Pods, and set the GOOGLE_APPLICATION_CREDENTIALS environment variable
- C. Create a new Google service account, and assign the appropriate IAM permissions
- D. Create a new JSON service account key for the Google service account, store the key in the secret management store for the CI/CD tool, and configure Terraform to use this key for authentication
- E. Assign the appropriate IAM permissions to the Google service account associated with the Compute Engine VM instances that run the Pods

Answer: A,C

Question: 175

[Applying Site Reliability Engineering (SRE) principles to a service]

You need to introduce postmortems into your organization during the holiday shopping season. You are expecting your web application to receive a large volume of traffic in a short period. You need to prepare your application for potential failures during the event. What should you do?

Choose 2 answers

- A. Monitor latency of your services for average percentile latency.
- B. Review your increased capacity requirements and plan for the required quota management.
- C. Create alerts in Cloud Monitoring for all common failures that your application experiences.
- D. Ensure that relevant system metrics are being captured with Cloud Monitoring and create alerts at levels of interest.
- E. Configure Anthos Service Mesh on the application to identify issues on the topology map.

Answer: B,D

Question: 176

[Building and implementing service monitoring strategies]

You are responsible for creating and modifying the Terraform templates that define your Infrastructure. Because two new engineers will also be working on the same code, you need to define a process and adopt a tool that will prevent you from overwriting each other's code. You also want to ensure that you capture all updates in the latest version.

What should you do?

- A. • Store your code in a Git-based version control system. • Establish a process that allows developers to merge their own changes at the end of each day. • Package and upload code to a versioned Cloud Storage bucket as the latest master version.
- B. • Store your code in a Git-based version control system. • Establish a process that includes code reviews by peers and unit testing to ensure integrity and functionality before integration of code. • Establish a process where the fully integrated code in the repository becomes the latest master version.
- C. • Store your code as text files in Google Drive in a defined folder structure that organizes the files. • At the end of each day, confirm that all changes have been captured in the files within the folder structure. • Rename the folder structure with a predefined naming convention that increments the version.
- D. • Store your code as text files in Google Drive in a defined folder structure that organizes the files. • At the end of each day, confirm that all changes have been captured in the files within the folder structure and create a new .zip archive with a predefined naming convention. • Upload the .zip archive to a versioned Cloud Storage bucket and accept it as the latest version.

Answer: B

Question: 177

[Building and implementing service monitoring strategies]

You created a Stackdriver chart for CPU utilization in a dashboard within your workspace project. You want to share the chart with your Site Reliability Engineering (SRE) team only. You want to ensure you follow the principle of least privilege. What should you do?

- A. Share the workspace Project ID with the SRE team. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
- B. Share the workspace Project ID with the SRE team. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.
- C. Click "Share chart by URL" and provide the URL to the SRE team. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
- D. Click "Share chart by URL" and provide the URL to the SRE team. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.

Answer: C

Question: 178

[Building and implementing CI/CD pipelines for a service]

As part of your company's initiative to shift left on security, the infoSec team is asking all teams to implement guard rails on all the Google Kubernetes Engine (GKE) clusters to only allow the deployment of trusted and approved images. You need to determine how to satisfy the InfoSec teams goal of shifting left on security. What should you do?

- A. Deploy Falco or Twistlock on GKE to monitor for vulnerabilities on your running Pods
- B. Configure Identity and Access Management (IAM) policies to create a least privilege model on your GKE clusters
- C. Use Binary Authorization to attest images during your CI CD pipeline
- D. Enable Container Analysis in Artifact Registry, and check for common vulnerabilities and exposures (CVEs) in your container images

Answer: C

Question: 179

[Managing service incidents]

Your company has a Google Cloud resource hierarchy with folders for production test and development. Your cyber security team needs to review your company's Google Cloud security posture to accelerate security issue identification and resolution. You need to centralize the logs generated by Google Cloud services from all projects only inside your

production folder to allow for alerting and near-real time analysis. What should you do?

- A. Enable the Workflows API and route all the logs to Cloud Logging
- B. Create a central Cloud Monitoring workspace and attach all related projects
- C. Create an aggregated log sink associated with the production folder that uses a Pub Sub topic as the destination
- D. Create an aggregated log sink associated with the production folder that uses a Cloud Logging bucket as the destination

Answer: D

Question: 180

[Optimizing service performance]

Your applications performance in Google Cloud has degraded since the last release You suspect that downstream dependencies might be causing some requests to take longer to complete You need to investigate the issue with your application to determine the cause What should you do?

- A. Configure Error Reporting in your application
- B. Configure Google Cloud Managed Service for Prometheus in your application
- C. Configure Cloud Profiler in your application
- D. Configure Cloud Trace in your application

Answer: D

Question: 181

[Building and implementing service monitoring strategies]

You are creating Cloud Logging sinks to export log entries from Cloud Logging to BigQuery for future analysis Your organization has a Google Cloud folder named Dev that contains development projects and a folder named Prod that contains production projects Log entries for development projects

must be exported to dev_dataset. and log entries for production projects must be exported to prod_dataset You need to minimize the number of log sinks created and you want to ensure that the log sinks apply to future projects What should you do?

- A. Create a single aggregated log sink at the organization level.
- B. Create a log sink in each project
- C. Create two aggregated log sinks at the organization level, and filter by project ID
- D. Create an aggregated log sink in the Dev and Prod folders

Answer: A

Question: 182

[Building and implementing CI/CD pipelines for a service]

You are configuring a CI pipeline. The build step for your CI pipeline integration testing requires access to APIs inside your private VPC network. Your security team requires that you do not expose API traffic publicly. You need to implement a solution that minimizes management overhead. What should you do?

- A. Use Cloud Build private pools to connect to the private VPC.
- B. Use Cloud Build to create a Compute Engine instance in the private VPC. Run the integration tests on the VM by using a startup script.
- C. Use Cloud Build as a pipeline runner. Configure a cross-region internal Application Load Balancer for API access.
- D. Use Cloud Build as a pipeline runner. Configure a global external Application Load Balancer with a Google Cloud Armor policy for API access.

Answer: A

Question: 183

[Building and implementing CI/CD pipelines for a service]

You recently created a Cloud Build pipeline for deploying Terraform code stored in a GitHub repository. You make Terraform code changes in short-lived branches and sometimes use tags during development. You tag releases with a semantic version when they are ready for deployment. You require your pipeline to apply the Terraform code whenever there is a new release, and you need to minimize operational overhead. What should you do?

- A. Create a build trigger with the * branch pattern.
- B. Create a build trigger with the \d+\.\d+\.\d* tag pattern.
- C. Create a build trigger with the .* tag pattern.
- D. Create a build trigger with the \d*\.\d+\.\d* branch pattern.

Answer: B

Question: 184

[Applying Site Reliability Engineering (SRE) principles to a service]

You have an application deployed to Cloud Run. A new version of the application has recently been deployed using the canary deployment strategy. Your Site Reliability Engineering (SRE) teammate informs you that an SLO has been exceeded for this application. You need to make the application healthy as quickly as possible. What should you do first?

- A. Configure traffic splitting to send 100% of the traffic to the latest revision.
- B. Configure traffic splitting to send 100% of the traffic to the previous revision.
- C. Create a new revision using the last known good version of the application.

D. Identify the cause of the latency by using Cloud Trace.

Answer: B

Question: 185

[Building and implementing service monitoring strategies]

You are responsible for the reliability of a custom-built, distributed file storage service that your company uses internally. This service handles thousands of file uploads and downloads daily. You need to define a service level indicator (SLI) to measure the reliability of your service usage and configure alerts to be notified of potential issues.

Which SLI should you use to measure the reliability of the service?

- A. Average request latency of API calls (e.g. get, put, list)
- B. Average size of objects stored in your service
- C. Ratio of successful API calls to the total number of attempted API calls
- D. Number of successful file uploads and downloads per minute

Answer: C

Question: 186

[Applying Site Reliability Engineering (SRE) principles to a service]

You are developing a Node.js utility on a workstation in Cloud Workstations by using Code OSS. The utility is a simple web page, and you have already confirmed that all necessary firewall rules are in place. You tested the application by starting it on port 3000 on your workstation in Cloud Workstations, but you need to be able to access the web page from your local machine. You need to

follow Google-recommended security practices. What should you do?

- A. Allow public IP addresses in the Cloud Workstations configuration.
- B. Use a browser running on a bastion host VM.
- C. Run the `gcloud compute start-iap-tunnel` command to the Cloud Workstations VM.
- D. Click the preview link in the Code OSS panel.

Answer: D

Question: 187

[Managing service incidents]

Your company runs an e-commerce business. The application responsible for payment processing has structured JSON logging with the following schema:

Capture and access of logs from the payment processing application is mandatory for operations, but the `jsonPayload.user_email` field contains personally identifiable information (PII). Your security team does not want the

entire engineering team to have access to PII. You need to stop exposing PII to the engineering team and restrict access to security team members only. What should you do?

- A. Apply a `jsonPayload.user_email` exclusion filter to the `_Default` bucket.
- B. Apply the conditional role binding `resource.name.extract(["locations/global/buckets/{bucket}/"]) == "_Default"` to the `_Default` bucket.
- C. Apply a `jsonPayload.user_email` restricted field to the `_Default` bucket. Grant the Log Field Accessor role to the security team members.
- D. Modify the application to toggle inclusion of `user_email` when the `log_user_email` environment variable is set to true. Restrict the engineering team members who can change the production environment variable by using the CODEOWNERS file.

Answer: C

Question: 188

[Applying Site Reliability Engineering (SRE) principles to a service]

You are designing a new multi-tenant Google Kubernetes Engine (GKE) cluster for a customer. Your customer is concerned with the risks associated with long-lived credentials use. The customer requires that each GKE workload has the minimum Identity and Access Management (IAM) permissions set following the principle of least privilege (PoLP). You need to design an IAM impersonation solution while following Google-recommended practices. What should you do?

- A. Create a Google service account. Create a Kubernetes service account in a Workload Identity-enabled cluster. Link the Google service account with the Kubernetes service account by using the `roles/iam.workloadidentityUser` role and `iam.gke.io/gcp-service-account` annotation. Map the Kubernetes service account to the workload. Repeat for each workload.
- B. Create a Google service account. Create a node pool, and set the Google service account as the default identity. Ensure that workloads can only run on the designated node pool by using node selectors, taints, and tolerations. Repeat for each workload.
- C. Create a Google service account. Create a service account key for the Google service account. Create a Kubernetes secret with a service account key. Ensure that workload mounts the secret and set the `GOOGLE_APPLICATION_CREDENTIALS` environment variable to point at the mount path. Repeat for each workload.
- D. Create a Google service account. Create a node pool without taints, and set the Google service account as the default identity. Grant IAM permissions to the Google service account.

Answer: A

Question: 189

[Building and implementing service monitoring strategies]

Your company stores a large volume of infrequently used data in Cloud Storage. The projects in your company's

CustomerService folder access Cloud Storage frequently, but store very little data. You want to enable Data Access audit logging across the company to identify data usage patterns. You need to exclude the CustomerService folder projects from Data Access audit logging. What should you do?

- A. Enable Data Access audit logging for Cloud Storage for all projects and folders, and configure exempted principals to include users of the CustomerService folder.
- B. Enable Data Access audit logging for Cloud Storage at the organization level, with no additional configuration.
- C. Enable Data Access audit logging for Cloud Storage at the organization level, and configure exempted principals to include users of the CustomerService folder.
- D. Enable Data Access audit logging for Cloud Storage for all projects and folders other than the CustomerService folder.

Answer: C

Question: 190

[Building and implementing service monitoring strategies]

Your organization is running multiple Google Kubernetes Engine (GKE) clusters in a project. You need to design a highly-available solution to collect and query both domain-specific workload metrics and GKE default metrics across all clusters, while minimizing operational overhead. What should you do?

- A. Use Prometheus Operator to install Prometheus in every cluster and scrape the metrics. Ensure that a Thanos sidecar is enabled on every Prometheus instance. Configure Thanos in the central cluster. Query the central Thanos instance.
- B. Use Prometheus Operator to install Prometheus in every cluster and scrape the metrics. Configure remote-write to one central Prometheus. Query the central Prometheus instance.
- C. Enable managed collection on every GKE cluster. Query the metrics in Cloud Monitoring.
- D. Enable managed collection on every GKE cluster. Query the metrics in BigQuery.

Answer: C

Question: 191

[Applying Site Reliability Engineering (SRE) principles to a service]

You work for a healthcare company and regulations require you to create all resources in a United States-based region. You attempted to create a secret in Secret Manager but received the following error message: The resource cannot be created in the global location due to an organization policy restriction. You need to resolve the error while remaining compliant with regulations. What should you do?

- A. Create the secret with an automatic replication policy.
- B. Create the secret with a user-managed replication policy.
- C. Remove the organization policy referenced in the error message.
- D. Add the global region to the organization policy referenced in the error message.

Answer: B

Question: 192

[Optimizing service performance]

You manage a retail website for your company. The website consists of several microservices running in a GKE Standard node pool with node autoscaling enabled. Each microservice has resource limits and a Horizontal Pod Autoscaler configured. During a busy period, you receive alerts for one of the microservices. When you check the Pods, half of them have the status OOMKilled, and the number of Pods is at the minimum autoscaling limit. You need to resolve the issue. What should you do?

- A. Increase the memory resource limit of the microservice.
- B. Increase the maximum number of nodes in the node pool.
- C. Increase the maximum replica limit of the Horizontal Pod Autoscaler.
- D. Update the node pool to use a machine type with more memory.

Answer: A

Question: 193

[Building and implementing CI/CD pipelines for a service]

Your company wants to implement a CD pipeline in Cloud Deploy for a web service deployed to GKE. The web service currently does not have any automated testing. The Quality Assurance team must manually verify any new releases of the web service before any production traffic is processed. You need to design the CD pipeline. What should you do?

- A. Create two pipeline stages, and use a canary deployment strategy.
- B. Create a single pipeline stage, and use a standard deployment strategy.
- C. Create a single pipeline stage, and use a canary deployment strategy.
- D. Create two pipeline stages, and use a standard deployment strategy.

Answer: D

Question: 194

[Optimizing service performance]

Your team has an application built by using a Dockerfile. The build is executed from Cloud Build, and the resulting artifacts are stored in Artifact Registry. Your team is reporting that builds are slow. You need to increase build speed, while following Google-recommended practices. What should you do?

- A. Use the --cache-from parameter, and point to Artifact Registry. Add the most frequently modified files to the later stages of the build process.

- B. Use the --cache-from parameter, and point to Artifact Registry. Add the most frequently modified files to the earlier stages of the build process.
- C. Cache the container layers of the build process to Cloud Storage. Add the most frequently modified files to the earlier stages of the build process.
- D. Cache the container layers of the build process to Cloud Storage. Add the most frequently modified files to the later stages of the build process.

Answer: A