



Navigation

Google Certified Professional Cloud DevOps Engineer Exam Prep

2 hours

★ 50 Questions

2.4 Minutes per Question


[Advanced \(/search?type=Practice Exam Challenge&difficulty=Advanced&categories=Google Cloud\)](#)

Go Back

Start Challenge

Question List

Show All Answers



← Go Back

Congratulations!

You passed this challenge on this attempt.

Expectations Report Card

DevOps Question Pool

88%

Exam Breakdown

DevOps Question Pool



1. Site Reliability Engineering maintains that 100% availability is not a good objective. Which of the following reasons is NOT one of the reasons?



A Users don't need 100% availability for a service to be acceptable.

B Less than 100% leaves room for new features.

C Achieving 100% availability is technically impossible

D The closer you get to 100%, the more expensive the process.

Correct Answer: C

Why is this correct?

While it is not easy – or the correct choice for all the other reasons listed here – achieving 100% availability for a specific period of time is doable and, thus, is not one of the reasons Site Reliability Engineers maintain that goal is recommended.

2. The team has recently set an SLO of 99.8% and an SLA of 99.9% for a particular service. Given this information, what is your error budget?



A More information is needed to calculate the error budget.

B The error budget would be 2%.

C The error budget would be .1%.

D The error budget would be .2%.

Correct Answer: D

Why is this correct?

The formula for calculating an error budget is $100\% - \text{SLO}$. In this scenario, this becomes $100\% - 99.8\% = .2\%$.

3. Which of these scenarios require installing and configuring custom logging agents on your GCP resources? Pick TWO options.



A Pod logs on a GKE cluster.

B Admin activity logs on a Compute Engine instance.

C Apache-access logs on a Compute Engine instance.

D View Windows Event Viewer logs on a Compute Engine instance.

Correct Answer: C

Why is this correct?

Third party application logs (like Apache) require the logging agent to be installed.

Correct Answer: D

Why is this correct?

OS logs (syslog and Windows Events) require the logging agent to be installed/configured.

4. In a production meeting, which of the following should NOT be discussed?



A Upcoming production changes.

B Outages.

C Personnel responsible for outages.

D Metrics.

Correct Answer: C

Why is this correct?

One of the pillars of Site Reliability Engineering is the embrace of a blameless culture.

5. You are creating a pipeline using Cloud Build to deploy a containerized application to Cloud Run. When Cloud Build runs, you receive an error that it does not have appropriate permissions to deploy to Cloud Run. What must you do to resolve this using Google's best practices?



A In Cloud Run permissions, assign the Cloud Build service account as a contributor in the Cloud Run project you are using.

B Assign the Project Owner role to the Cloud Build service account.

C Assign the Project Owner role to the end user deploying the Cloud Build pipeline.

D Assign the Cloud Run Admin role to the Cloud Build service account.

Correct Answer: D

Why is this correct?

The Cloud Build service account needs permission to manage Cloud Run.

6. You are designing a deployment pipeline to deploy containerized images to GKE. You need to monitor and be informed of known vulnerabilities in built container images. What should you do?



A Enable Vulnerability Scanning in Cloud Build to scan added images for known vulnerabilities during the build process.

B Enable vulnerability scanning in Container Registry to scan added images for known vulnerabilities.

C Implement Binary Authorization to scan container images for vulnerabilities before deploying to GKE.

D Enable Vulnerability Scanning in Cloud Source Repositories to scan added images for known vulnerabilities during the build process.

Correct Answer: B

Why is this correct?

The Vulnerability Scanning feature is enabled and managed in the Container Registry service.

7. According to SRE best practices, which of the following is NOT a true statement?



A There are no consequences for missing an SLO, only for missing an SLA.

B One consequence of missing an SLO is slowing down the rate of change.

C One consequence of missing an SLO is directing more engineering efforts at reducing risk and improving reliability.

D An SLO should represent a level of the desired user experience.

Correct Answer: A

Why is this correct?

SLOs and SLAs should both have consequences to be meaningful and incentivizing.

8. In a production meeting, which paging events should be considered?



A Events that were paged but should not have been.

B All of these options.

C Events that were paged.

D Events that were not paged, but should have been.

Correct Answer: B

Why is this correct?

In a production meeting, the full range of pages over the relevant time period should be discussed.

9. You are using Cloud Build as part of a larger CI/CD pipeline. When Cloud Build finishes building a new container image, your continuous delivery tool needs to be notified. How is this interaction configured?



A Integrate your continuous delivery tool with the Cloud Build CD connector.

B In the Cloud Build configuration file, configure a webhook to notify the continuous delivery application of the new image.

C Subscribe the continuous delivery tool to the default Pub/Sub topic that Cloud Build uses to notify of new completed builds.

D Enable GKE integration with Cloud Build for GKE to be notified of new images.

Correct Answer: C

Why is this correct?

This is best practice. You can also configure custom Pub/Sub topics as an alternative to utilizing the default one used by Cloud Build.

10. You are planning the deployment of an application to Kubernetes Engine. You need to collect detailed metrics across all aspects of the deployment including node-utilization and container metrics. What do you need to do to enable the collection of these detailed metrics?



A Install and configure the Monitoring Agent on your GKE cluster.

B Enable Cloud Operations for GKE in your cluster settings.

C Enable container monitoring into your build pipeline.

D Install and configure the Logging Agent on your GKE cluster.

Correct Answer: B

Why is this correct?

Cloud Operations for GKE enables detailed metrics and logging across all aspects of your GKE cluster. Be aware that this feature is automatically enabled on newer clusters.

11. Which of the following sets-out in advance what happens (the consequences) when availability of an application fails?



A SLO.

B SLI.

C PRR.

D SLA.

Correct Answer: D

Why is this correct?

A Service Level Agreement is a document that describes the consequences should agreed-upon service levels of availability are not met.

INCORRECT

12. You are calculating an availability SLI for an eCommerce website where, over a month, The HTTP GET requests for `/checkout_service/response_counts` are divided, as follows:
2xx status: 900,000 3xx status: 40,000 4xx status: 50,000 5xx status: 10,000
What is the SLI?



A 90.0%.

B 99.0%.

C 98.9%

D 99.9%.

Your Answer: A

Why is this incorrect?

This calculation does not exclude invalid events such as 3xx and 4xx responses.

Correct Answer: C

Why is this correct?

The formula for calculating an SLI is (Good Events) / (Valid Events) X 100 – which, because we exclude the invalid 3xx and 4xx responses, would equal 98.9%.

INCORRECT

13. You are planning a Monitoring workspace(s) structure that will include multiple projects. You need to monitor the following projects which are managed by separate teams: app1-dev app1-test app1-qa app1-prod Team members who have access to the app1-dev and app1-test projects need to view performance data in those projects, but should not have access to view performance data in app1-qa or app1-prod. Team members who have access to the app1-qa and app1-prod projects need to view performance data in those projects, but should not have access to view performance data in app1-dev or app1-test. How should you structure your Cloud Monitoring workspaces to fulfill this requirement using the principle of least privilege?  

A Create two Cloud Monitoring workspaces, each one in a separate project called 'workspace-1' and 'workspace-2'. Assign projects app1-dev and app1-test to the monitoring workspace in workspace-1. Assign projects app1-qa and app1-prod to the monitoring workspace in workspace-2. Assign team members the Monitoring Editor role to the workspace projects.

B Create a single Cloud Monitoring workspace in a dedicated project, and assign all projects to the workspace. Assign team members in the associated projects the Monitoring Viewer role to view only their own projects in the workspace.

C Create two Cloud Monitoring workspaces, each one in a separate project called 'workspace-1' and 'workspace-2'. Assign projects app1-dev and app1-test to the monitoring workspace in workspace-1. Assign projects app1-qa and app1-prod to the monitoring workspace in workspace-2. Assign team members the Monitoring Viewer role to the workspace projects.

D Create two Cloud Monitoring workspaces named 'workspace-1' and 'workspace-2' in a single project. Assign projects app1-dev and app1-test to workspace-1. Assign projects app1-qa and app1-prod to workspace-2. Assign team members the Monitoring Viewer role to the individual workspaces.

Your Answer: D



Why is this incorrect?

You can only create a single workspace per project, which takes on the same name as the project ID. Therefore, this is not a valid solution.

Correct Answer: C

Why is this correct?

Team members who have access to a workspace can view performance data for all projects associated with that workspace. Therefore, we need to separate monitored projects into two separate workspaces based on access scopes. The Monitoring Viewer role will limit users to view but not edit workspace data.

14. You are planning an automated CI/CD pipeline to use for developing your GKE application. Your pipeline needs to support the implementation of separate blue/green deployments to easily switch from one version to another, as well as support the ability to easily roll back to a prior version of your application. This process must use automation whenever possible. How should you configure this pipeline?  

A Use Cloud Build as a continuous integration tool to build and test your code. Update and apply your Kubernetes YAML files to deploy different versions of each application.

B Use Cloud Build as a continuous integration tool to build and test your code. Use Spinnaker as a continuous delivery tool to implement separate staging and production (or blue/green) pipelines.

C Use Cloud Build as both a continuous integration and delivery tool to manage application deployments.



D Use Cloud Source Repositories as a continuous integration tool to build and test your code. Use Spinnaker as a continuous delivery tool to implement separate staging and production (or blue/green) pipelines.

Correct Answer: B

Why is this correct?

Spinnaker (as well as Jenkins) are fully featured continuous delivery tools that support deployment methods such as blue/green and canary deployments.

INCORRECT

15. Your team is developing a new application using Google Cloud Vision API to categorize labels from wine bottles, tying their recognition to a database of details including recent prices and reviews. People will be able to take a picture of a label and very quickly retrieve the information without searching for the wine by name. You've been tasked with developing the SLIs for the app. Which of the following is the best selection of SLIs to use?  

A Availability, Latency, and Quality.

B Freshness, Correctness, and Throughput.

C Latency, Correctness, and Durability.

D Availability, Freshness, and Quality.

Your Answer: A



Why is this incorrect?

Availability, latency, and quality are considered better for more request-driven types of applications.

Correct Answer: B

Why is this correct?

Freshness reflects the most recently updated data; correctness is the proportion of accurate records returned; and throughput is the proportion of time when the data is faster than a defined threshold. All three are important to a data-driven app.

16. Which of the following IAM roles grants the SMALLEST scope of rights for a service account to write logs to Cloud Logging?  

A Logs Writer

B Logs Configuration Writer

C Logging Admin

D Private Logs Viewer

Correct Answer: A

Why is this correct?

The Logs Writer role allows a service account to write logs to Cloud Logging without giving it read permissions.

17. Which of the following security logs types are NOT enabled by default?



A Admin Activity

B System Event

C Data Access

D Access Transparency

Correct Answer: C

Why is this correct?

Data Access logs are NOT enabled by default, must be manually enabled and can be disabled.

18. What is the preferred method for the Cloud Build service to inform a separate CD tool that a new artifact is available for provisioning?



A Cloud Build publishes a message to a Pub/Sub topic when an artifact is completed. The CD tool is subscribed to the same topic to be notified of completion.

B Cloud Build sends completion messages to a Cloud Storage bucket. The CD tool will watch the Cloud Storage bucket for notification of completion.

C Cloud Build sends an HTTP/POST request to the CD tool.

D Cloud Build notifies the CD tool via an HTTP webhook.

Correct Answer: A

Why is this correct?

Cloud Build uses a default Pub/Sub topic for notifications, or you can also configure a custom topic as well.

19. You're the engineer on duty and you receive a text message that a data center has gone offline shortly after a new release of your application went live. Following Google Cloud SRE best practices, what should you do?



A Send a notice to the public relations team so they can let the public know that you are aware of the situation and are working diligently to fix it.

B Send an "all hands on-deck" alert to your fellow team members.

C Rollback your latest release to correct the problem.

D Designate a development team member as the incident commander.

Correct Answer: D

Why is this correct?

Designating an incident commander – or taking that role yourself - is the first step in a best practice for incident response.

20. While going over the latest monitoring metrics, you notice an increase in latency for your application during a specific time period in the previous week. There were no deployments during that period so you suspect the issue may have been a problem with the Google infrastructure. What's the most efficient way to proceed?



- A Run a diagnostic test on your app to review current latency figures.
- B Contact Google Cloud support by email.
- C Check your admin audit logs for the app to see if Google made any administrative changes that could have affected latency.
- D Review the APIs and Services Dashboard to see if there are any corresponding increases in latency for the given period.**

Correct Answer: D

Why is this correct?

Google Cloud's APIs and Services Dashboard gives you a "transparent SLI" that shows you how the various APIs and services are performing, complete with charts detailing Median Latency over the past 30 days.

21. You are creating a scaling Nginx web application running on GCE instances on a managed instance group. For troubleshooting purposes, you need to collect custom logs from the instances' /var/logs directory and export them to Cloud Logging. These custom logs fall outside of typical third-party application use cases. You also need to collect typical third-party logs from Nginx as well. What should you do to capture these logs?



- A In your managed instance group template, configure a startup script to automatically install the Logging Agent and catch-all configuration to capture standard Nginx logs.
- B On each instance in the instance group, install the Logging Agent and catch-all configuration to capture custom logs.
- C On each instance in the instance group, install the Logging Agent and catch-all configuration to capture standard Nginx logs. Configure the logging agent to also capture custom logs.
- D In your managed instance group template, configure a startup script to automatically install the Logging Agent and catch-all configuration to capture standard Nginx logs. In the same startup script, configure the logging agent to also capture custom logs.**

Correct Answer: D

Why is this correct?

We want to automatically install and configure the agent to capture default and custom logs. This is the correct answer.

22. According to Google Cloud recommended best practices, when can an application or service move from limited availability to general availability?



A When the Production Readiness Review has been passed.

B When the development team lead indicates that they have had no incidents or issues for the past agreed-upon period.

C When the GA date designated by management is reached.

D When the SRE team lead and the development team lead agree that development is complete.

Correct Answer: A

Why is this correct?

The Production Readiness Review (PRR) is a process that identifies the reliability needs of a service based on its specific details. Through a PRR, SREs seek to apply what they've learned and experienced to ensure the reliability of a service operating in production.

23. Your organization uses Anthos to manage Kubernetes clusters in multiple environments. You are planning the deployment of an application that will run on GKE. Due to compliance requirements, we need to control which images are allowed to run on our application's cluster. No 'unapproved' images can be allowed to run on our cluster. What is the preferred mechanism for controlling which approved images can run on your cluster?



A Implement Binary Authorization to digitally sign approved images before they are deployed onto your locked-down cluster.

B Use a secrets management system to manage which images are allowed on your cluster.

C Implement Binary Authorization to digitally sign approved images after they are deployed onto your locked-down cluster.

D Enable image filtering in GKE, and specify the name of approved image ID's. All image IDs not on the list will be rejected.

Correct Answer: A

Why is this correct?

Binary Authorization uses a built attester image to digitally sign approved images before they are deployed to your GKE cluster.

24. Which of the following is NOT considered a best practice when it comes to incident response management?



A Prioritize damage and restore service.

B Incident commander must sign off on operation team fixes.

C Develop and document incident response procedures.

D Trust team members in specified roles.

Correct Answer: B

Why is this correct?

All team members, including the operation team, should be autonomous and allowed to make the decisions and take the actions they think best.

25. Which of these scenarios require installing and configuring custom monitoring agents on your GCP resources? Pick TWO options.



- A View detailed metrics on your GKE cluster.
- B View average CPU utilization of Compute Engine instances.
- C View Apache web server access metrics on Compute Engine.**
- D View memory usage metrics on Compute Engine instances.**

Correct Answer: C

Why is this correct?

Third party metrics (like Apache) require the monitoring agent to be installed.

Correct Answer: D

Why is this correct?

Memory usage requires the monitoring agent to be installed.

26. How can you monitor metrics from multiple projects at the same time?



- A Create a metrics export in each project to a single BigQuery dataset.
- B Create a single workspace in a project, and add each monitored project to that workspace.**
- C A project's metrics can only be viewed from within the project they come from.
- D Create a project for the sole purpose of monitoring. Create a workspace for each monitored project in this project.

Correct Answer: B

Why is this correct?

You can monitor multiple projects from a single workspace. You can only have a single workspace per project.

INCORRECT

27. You are planning on deploying Jenkins on GCP to manage artifact builds. What is GCP's recommendation to deploy Jenkins to reduce operational overhead?



- A Deploy Jenkins with a single master with agents operating in a managed instance group.**
- B Deploy Jenkins on-premises, using assets hosted on Google Cloud Storage to deploy to GCP.

C Deploy Jenkins to App Engine.

D Deploy Jenkins to GKE.

Your Answer: A

Why is this incorrect?

This option certainly exists and is documented by Google, however running Jenkins on GKE offers lower overhead and better performance. The instance group route is for those learning to transition from on-premises environments.

Correct Answer: D

Why is this correct?

Running Jenkins as a containerized application on GKE requires substantially less operational overhead.

INCORRECT

28. Your organization uses separate teams for developing applications on GKE in separate projects. Each team member should only access the cluster in their own project. You want to configure Cloud Build to automatically process builds for each team. How should you configure Cloud Build to enable each team to process builds for their cluster without exposing additional access?



A Create a Cloud Build configuration inside of a single 'master' project. Grant the Kubernetes Engine Developer role from each team's project to the Cloud Build service account.

B Use the profiles option built into Cloud Build to limit access to each project.

C Create a Cloud Build configuration inside of a single 'master' project. Create a separate service account for each team's project to be associated with Cloud Build. Grant the Kubernetes Engine Developer role from each team's project to their specific Cloud Build service account.

D Create a Cloud Build configuration in each team's project. In each project, grant the Cloud Build service account the Kubernetes Engine Developer role.

Your Answer: A

Why is this incorrect?

While technically possible, this will give each team build access to every other team's cluster, which is too much access. We need to configure Cloud Build in each separate project, scoped for only that project.

Correct Answer: D

Why is this correct?

Cloud Build is designed to run jobs inside of its own project. Cloud Build needs access to the GKE cluster in order to complete builds

29. You are troubleshooting connectivity issues to a Compute Engine instance on your network. You suspect that traffic is getting mistakenly blocked by misconfigured firewall rules. What troubleshooting method can you use to accurately determine this?



A Enable VPC Logging on your VPC network to view denied connection attempts by network port. Export logs to BigQuery for analysis.

B Enable Firewall Logging on the firewall rules you suspect may be blocking traffic. Export logs to BigQuery for analysis.

C Enable Firewall Logging on your affected subnet. Export logs to BigQuery for analysis.

D Enable VPC Logging on the affected firewall rule.

Correct Answer: B

Why is this correct?

Enabling Firewall logs will give insight as to which connections are allowed and blocked.

30. Your web application using a managed instance group has been generating a large number of alerts determined to be 'false positives' that do not require immediate action. These alerts are generated when average instance CPU usage is above 80% for 1 minute. You need to modify your alerting strategy to alert you when your instances are legitimately over-utilized WITHOUT impacting detection time. What step can you take?



A Adjust your alert condition to only alert when CPU utilization is over 95% for a minute vs. 80%.

B Increase the duration of the alert condition to 5 minutes. Leave the metric aggregation to 1 minute.

C Adjust your CPU utilization metric to measure average CPU utilization over a 20-minute period, and alert when average CPU utilization over a 20-minute period aggregation exceeds 95%.

D Adjust your CPU utilization metric to measure average CPU utilization over a 20-minute period, and alert when average CPU utilization over a 20-minute period aggregation exceeds 80%.

Correct Answer: B

Why is this correct?

Increasing the duration of an alert condition increases precision while not too negatively impacting detection time.

31. Your globally-based GKE application utilizes a secret to protect sensitive information. The application uses multiple GKE clusters around the globe for fast access to users on multiple continents. How should you structure your secrets management for your application according to Google's best practices?



A Store the secret in the Cloud DLP service. Configure the application to reference the secret stored in Cloud DLP.

B Store the secret in the Google Secrets Manager. Configure the application to reference the secret stored in Cloud KMS.

C Create an encoded secret on each cluster. Use a cron job to synchronize secret changes on each cluster.

D Include the secret in the application source code to be built into the application upon runtime.

Correct Answer: B

Why is this correct?

Google Secrets Manager acts as a central repository for secrets, and is the ideal solution for this scenario.

32. Your organization is required to satisfy a number of compliance requirements. One of those requirements is to store Admin Activity logs for 10 years to meet audit requirements. What should you



do to meet this requirement?

- A Stream the admin activity logs into a Pub/Sub topic, which will then connect to a Cloud Storage bucket for retention.
- B Set a custom retention period for Admin Activity logs to 10 years.
- C Export the logs via a logs-based metric, and set this metric for a 10 year retention period.
- D Create an export sink to capture all admin activity logs via a filter, and export them to Cloud Storage. Set a storage lifecycle policy in Cloud Storage to retain objects for 10 years.**

Correct Answer: D

Why is this correct?

Export logs out of Cloud Logging to Cloud Storage allows you to keep exported logs for however long you wish. Within Cloud Logging, Admin Activity logs will be retained for 400 days, then be deleted. This is not configurable.

33. Your new app is deployed on-premises, around the world at various international locations for internal use, and on the Google Cloud Platform for external use. To assure most efficient use of resources, all traffic is routed through Google Cloud load balancers. Your SLIs need to gather latency and availability data across the board. What should you do?



- A Gather the data from the load balancers.**
- B Gather the data only from the GCP VMs as they represent your most important users.
- C Gather and process data together from the on-prem VMs and the GCP VMs to create a single bucketed SLI.
- D Gather and process data separately from the on-prem VMs and the GCP VMs to create separate SLIs which can then be aggregated.

Correct Answer: A

Why is this correct?

Google Cloud recommends gathering SLI data from load balancers as they are closer to the user and it is the most efficient method.

34. What tool should you use to analyze a sample of network traffic on a specific subnet on your VPC?





- A Logging agents
- B Firewall Logs
- C VPC Flow Logs**
- D Monitoring Agents

Correct Answer: C

Why is this correct?

VPC Flow Logs are enabled per subnet, and log a sample of all network traffic on that subnet.



35. You work for a global enterprise and an incident has occurred causes massive disruption to your team's application. You're the Incident Commander. The operations team has tried numerous solutions, but none have resolved the problem. It's nearing the end of your workday. What should you do?  

- A Designate a different operations team in a different time zone, so they can continue working on the issue.
- B Authorize overtime for the operations team so they can continue working on the problem.
- C Hand off responsibility to another incident commander in a different time zone, so he/she can continue working on the issue.**
- D Let corporate know so they can inform all the stakeholders that the issue will continue for another day.

Correct Answer: C

Why is this correct?

Best practices dictate that each incident have a single incident commander at all times and, if the original incident commander is unable to continue, he or she should hand off responsibility to a new one.

36. You've been asked to set up error reporting notification for a consultant on your team's project. Which of the following steps do you need to take? (Choose two from the available options.)  

- A Give the consultant the `errorreporting.viewer` role.
- B Give the consultant a custom role with `cloudnotifications.activities.list` permission.**
- C Enable error reporting notification for the project.**
- D Add the consultant's email address to the Google Cloud Error Reporting user interface for the project.

Correct Answer: B



Why is this correct?

Error reporting notifications are available to users with Project Owner, Project Editor, Project Viewer role or a custom role with the `cloudnotifications.activities.list` permission.

Correct Answer: C

Why is this correct?

Error reporting must be enabled for each project independently.

37. You've been tasked to set up the error budget burn rate alerting policies for your team's new app. You want to be sure to follow Google Cloud SRE best practices. What should you do?  



- A Set an alert to notify engineers whenever the established SLOs are exceeded.
- B Set an alert to notify engineers whenever the number of errors occurring over either a 12 hour period or a 5 minute period could exhaust the established error budget before the compliance period has expired.**

- C Set an alert to notify engineers whenever the number of errors occurring over a 12 hour period could exhaust the established error budget before the compliance period has expired.
- D Set an alert to notify engineers whenever the number of errors occurring over a 5 minute period could exhaust the established error budget before the compliance period has expired.

Correct Answer: B

Why is this correct?

Alerts that focus on the burn rate of an error budget require watching both the slow-burn (over an extended period, like 12 hours) and fast-burn (over a shorter period, such as 5 minutes) scenarios.



38. According to Google Cloud SRE best practices, which of the following is NOT an element of an incident postmortem.  

- A A recreation of the timeline.
- B Identify the team member(s) responsible for the incident and detail subsequent actions taken.**
- C Generate a report of the incident.
- D Gathering and documenting the incident metadata.

Correct Answer: B

Why is this correct?

One of the primary tenets of Site Reliability Engineering is the blameless culture regarding incidents. A blameless culture means that no one is at fault, no one is shamed, no one is fired, and everyone learns.



39. Which of the following is NOT a characteristic of white-box monitoring?  

- A The focus is on predicting problems.
- B Best for detecting imminent issues.
- C Includes metrics exposed by the internals of the system.
- D Used to test externally visible behavior as the user would see it.**

Correct Answer: D

Why is this correct?

Testing externally visible behavior is known as black-box monitoring.

40. Your application is hosted on a Linux Compute Engine instance. Your instance periodically crashes, leading you to suspect a fault in the OS. You want to query syslog entries in Cloud Logging. What do you need to do to achieve this?  

- A** Install the Logging Agent and 'catch-all' configuration on the affected instance. New syslog entries will automatically appear in Cloud Logging for analysis.
- B Install the Logging Agent and Linux OS configuration. New syslog entries will automatically appear in Cloud Logging for analysis.
- C Your GCE instance will automatically send OS syslog entries to Cloud Logging without additional configuration.
- D Enable the Cloud Operations for GCE feature on your instance to collect OS logs.

Correct Answer: A

Why is this correct?

This is correct. Installing the Logging Agent and catch-all configuration will automatically capture logs from syslog from the Linux OS.

41. How can you view a subset of logs from multiple projects in a folder at the same time?



- A Create a workspace, and add the necessary projects to that workspace to view logs in a single location.
- B Within the Cloud Logging view, add each project that needs to be viewed.
- C Access the Cloud Logging Logs Viewer feature in the folder, and view all project's logs from there.
- D** Create a logging export sink which includes the folder (as well as children objects) and the logging filter to be defined. Export these logs to another service such as BigQuery.

Correct Answer: D

Why is this correct?

This is correct. You cannot view logs in Cloud Logging in multiple projects, however you can export logs from multiple projects to a single source using a command line-level export.

42. When using Spinnaker as a continuous delivery tool, how does Spinnaker control traffic to different versions of your Kubernetes application?



- A Deploy a new service for each new version of a production deployment.
- B Updates existing Kubernetes Deployments for each new version.
- C** Updates the existing production Replica Set for each new version deployed.
- D Creates a new Replica Set for each new version deployed.

Correct Answer: C

Why is this correct?

Spinnaker updates existing staging and production Replica Sets with each new version of a deployed application.

43. Cloud Profiler supports which of the following computing systems?



A Compute Engine.

B Computing platforms outside of Google Cloud.

C All of these options.

D Kubernetes Engine.

Correct Answer: C

Why is this correct?

All of the listed options are supported.

INCORRECT

44. You've been assigned to figure out why your application appears to be degrading and it is in danger of missing its latency SLO. Which of the following Google Cloud services is best to use?



A Cloud Debugger.

B Cloud Monitoring.

C Cloud Logging.

D Cloud Profiler.

Your Answer: B

Why is this incorrect?

Cloud Monitoring collects metrics, events, and metadata from Google Cloud, Amazon Web Services, hosted uptime probes, and application instrumentation, but is not the best choice for tracking latency data.

Correct Answer: D

Why is this correct?

Cloud Profiler is designed to examine latency within applications and specifically reveal latency figures over time.

45. You manage a user-facing web application. What deployment model do you use to deploy a new, production version of an application to a small subset of users that can be rolled back in case of errors?



A Canary update

B Red/Black model

C Blue/Green model

D Test version

Correct Answer: A

Why is this correct?

Canary updates are deployed to a live production environment, but to only a small percentage of users alongside the previous version. If errors occur, the canary deployment is then rolled back.

46. Your team is transitioning to establishing Site Reliability Engineering standards for their new app and you've been assigned the job of recommending the initial SLIs to use. You want to keep to Google Cloud best practices. Which of the following do you NOT recommend? (Choose two from the following.)



A Latency.

B Error.

C Traffic.

D Freshness.

E Reliability.

F Saturation.

Correct Answer: D

Why is this correct?

Freshness is important for data-related apps, but not considered one of Google's Four Golden signals and, thus, should not be recommended.

Correct Answer: E

Why is this correct?

Reliability is not generally considered a service level indicator, and is definitely not considered one of Google's Four Golden signals and, thus, should not be recommended.

47. Select TWO recommended methods of improving Cloud Build performance from the provided list.



A Use lighter images such as Alpine Linux when possible.

B Exclude unnecessary files from builds using a .gcloudignore file.

C Compress the code image before building.

D Exclude unnecessary files from builds using a .gitignore file.

Correct Answer: A

Why is this correct?

Lighter images are smaller and take less time to build.

Correct Answer: B

Why is this correct?

A .gcloudignore file excludes files that don't need to be part of the build.

48. Your company just hired a new head of production who is very gung-ho. After examining the last quarter's up-time for your latest mobile app, he has stated that he wants to out-perform their primary competition, and shoot for 100% availability. What should you do?



- A** Inform the head of production that 100% availability is not a realistic goal.
- B Double the number of developers and revamp the product to insure the 100% availability.
- C Monitor the competition's latency, and send SMS alerts to marketing whenever their product's response time is over 100 ms.
- D Install a monitoring agent on all VMs, with SMS alerts sent to on-call engineers any time the service response is under the 100.00% latency SLO.

Correct Answer: A

Why is this correct?

A principal tenet of Site Reliability Engineering is that 100% availability is too expensive a goal; and that knowing limitations is actually a benefit which will allow for an error budget, and subsequent windows of opportunity to launch new features.

49. Which of the following is not a characteristic of 'toil', as defined by Site Reliability Engineering?



- A Repetitive.
- B Reactive and interrupt-driven.
- C Manual.
- D** Scales exponentially as service grows.

Correct Answer: D

Why is this correct?

Toil tends to scale linearly as the service grows.

50. You've been named the head of a committee to establish an incident management policy that would incorporate various team members. Which of the following roles is NOT essential to an incident response team?



- A** A social media team, responsible for managing expectations and informing the public.
- B One or more people designated with the actual responses (such as code fixes) to the incident.
- C The public face of the incident, responsible for issuing updates to stakeholders.
- D Someone who is in charge of the incident and designates responsibility for other roles.

Correct Answer: A

Why is this correct?

Why is this correct:

While there very well might be a social media team involved at some point, communications regarding the incident response should be handled by the Communications Lead.
