[28/July/2018 Updated Real Professional Cloud Architect Exam Questions -- All People Need To Learn For Not Failing Exam (Part B)

New Updated Professional Cloud Architect Exam Questions from PassLeader Professional Cloud Architect PDF dumps! Welcome to download the newest PassLeader Professional Cloud Architect VCE

dumps: https://www.passleader.com/professional-cloud-architect.html (82 Q&As) Keywords: Professional Cloud Architect exam dumps, Professional Cloud Architect exam questions, Professional Cloud Architect VCE dumps, Professional Cloud Architect PDF dumps, Professional Cloud Architect practice tests, Professional Cloud Architect study guide, Professional Cloud Architect braindumps, Google Cloud Certifications: Professional Cloud Architect Exam P.S. New Professional Cloud Architect dumps PDF: https://drive.google.com/open?id=19jt3GbCmVz-pmGbZv8zjAu0NH7423IQ2 P.S. New Professional Data Engineer dumps PDF: https://drive.google.com/open?id=1m882ngsiRO1BOHineV4IQUv9jgF5Lpue NEW QUESTION 16 The application reliability team at your company has added a debug feature to their backend service to send all server events to Google Cloud Storage for eventual analysis. The event records are at least 50 KB and at most 15 MB and are expected to peak at 3,000 events per second. You want to minimize data loss. Which process should you implement? A. Append metadata to file body. Compress individual files. Name files with serverName-Timestamp. Create a new bucket if bucket is older than 1 hour and save individual files to the new bucket. Otherwise, save files to existing bucket. B. Batch every 10,000 events with a single manifest file for metadata. Compress event files and manifest file into a single archive file. Name files using serverName-EventSequence. Create a new bucket if bucket is older than 1 day and save the single archive file to the new bucket. Otherwise, save the single archive file to existing bucket. C. Compress individual files. Name files with serverName-EventSequence. Save files to one bucket. Set custom metadata headers for each object after saving. D. Append metadata to file body. Compress individual files. Name files with a random prefix pattern. Save files to one bucket. Answer: A NEW QUESTION 17 A lead software engineer tells you that his new application design uses websockets and HTTP sessions that are not distributed across the web servers. You want to help him ensure his application will run property on Google Cloud Platform. What should you do? A. Help the engineer to convert his websocket code to use HTTP streaming. B. Review the encryption requirements for websocket connections with the security team. C. Meet with the cloud operations team and the engineer to discuss load balancer options. D. : : Help the engineer redesign the application to use a distributed user session service that does not rely on websockets and HTTP sessions. Answer: C NEW QUESTION 18 Your solution is producing performance bugs in production that you did not see in staging and test environments. You want to adjust your test and deployment procedures to avoid this problem in the future. What should you do? A. Deploy fewer changes to production. B. Deploy smaller changes to production. C. Increase the load on your test and staging environments. D. Deploy changes to a small subset of users before rolling out to production. Answer: C NEW QUESTION 19 Your company has decided to make a major revision of their API in order to create better experiences for their developers. They need to keep the old version of the API available and deployable, while allowing new customers and testers to try out the new API. They want to keep the same SSL and DNS records in place to serve both APIs. What should they do? A. Configure a new load balancer for the new version of the API. B. Reconfigure old clients to use a new endpoint for the new API. C. Have the old API forward traffic to the new API based on the path. D. Use separate backend pools for each API path behind the load balancer. Answer: B Explanation: https://cloud.google.com/endpoints/docs/openapi/lifecycle-management NEW QUESTION 20 A small number of API requests to your microservices-based application take a very long time. You know that each request to the API can traverse many services. You want to know which service takes the longest in those cases. What should you do? A. Set timeouts on your application so that you can fail requests faster. B. Send custom metrics for each of your requests to Stackdriver Monitoring. C. Use Stackdriver Monitoring to look for insights that show when your API latencies are high. D. Instrument your application with Stackdnver Trace in order to break down the request latencies at each microservice. Answer: D Explanation: https://cloud.google.com/trace/docs/overview NEW QUESTION 21 During a high traffic portion of the day, one of your relational databases crashes, but the replica is never promoted to a master. You want to avoid this in the future. What should you do? A. Use a different database. B. Choose larger instances for your database. C. Create snapshots of your database more regularly. D. Implement routinely scheduled failovers of your databases. Answer: C Explanation: https://cloud.google.com/solutions/disaster-recovery-cookbook NEW QUESTION 22 You need to reduce the number of

unplanned rollbacks of erroneous production deployments in your company's web hosting platform. Improvement to the QA/Test processes accomplished an 80% reduction. Which additional two approaches can you take to further reduce the rollbacks? (Choose two.) A. Introduce a green-blue deployment model. B. Replace the QA environment with canary releases. C. Fragment the monolithic platform into microservices. D. Reduce the platform's dependency on relational database systems. E. Replace the platform's relational database systems with a NoSOL database. Answer: AC NEW OUESTION 23 An application development team believes their current logging tool will not meet their needs for their new cloud-based product. They want a better tool to capture errors and help them analyze their historical log data. You want to help them find a solution that meets their needs, what should you do? A. Direct them to download and install the Google StackDriver logging agent. B. Send them a list of online resources about logging best practices. C. Help them define their requirements and assess viable logging tools. D. Help them upgrade their current tool to take advantage of any new features. Answer: A NEW QUESTION 24 Your company plans to migrate a multi-petabyte data set to the cloud. The data set must be available 24hrs a day. Your business analysts have experience only with using a SQL interface. How should you store the data to optimize it for ease of analysis? A. Load data into Google BigQuery. B. Insert data into Google Cloud SQL. C. Put flat files into Google Cloud Storage. D. Stream data into Google Cloud Datastore. Answer: A Explanation: Google Big Query is for multi peta byte storage, HA(High availability) which means 24 hours, SQL interface.

https://medium.com/google-cloud/the-12-components-of-google-bigquery-c2b49829a7c7

<u>https://cloud.google.com/solutions/bigquery-data-warehouse https://cloud.google.com/bigquery/</u> NEW QUESTION 25 Your application needs to process credit card transactions. You want the smallest scope of Payment Card Industry (PCI) compliance without compromising the ability to analyze transactional data and trends relating to which payment methods are used. How should you design your architecture? A. Create a tokenizer service and store only tokenized data.

B. Create separate projects that only process credit card data. C. Create separate subnetworks and isolate the components that process credit card data. D. Streamline the audit discovery phase by labeling all of the virtual machines (VMs) that process PCI data. E. Enable Logging export to Google BigQuery and use ACLs and views to scope the data shared with the auditor. Answer: E Explanation: The proper model for exporting credit card processing data is to forward from a squid proxy to Stackdriver Logging, and export from Stackdriver Logging into BigQuery. <u>https://cloud.google.com/solutions/pci-dss</u> NEW QUESTION 26 You have been asked to select the storage system for the click-data of your company's large portfolio of websites. This data is streamed in from a custom website analytics package at a typical rate of 6,000 clicks per minute, with bursts of up to 8,500 clicks per second. It must been stored for future analysis by your data science and user experience teams. Which storage infrastructure should you choose? A.

Cloud Storage D. Google cloud Datastore Answer: C Explanation:

https://cloud.google.com/solutions/data-analytics-partner-ecosystem

https://zulily-tech.com/2015/08/10/leveraging-google-cloud-dataflow-for-clickstream-processing/ NEW QUESTION 27 Your customer is receiving reports that their recently updated Google App Engine application is taking approximately 30 seconds to load for some of their users. This behavior was not reported before the update. What strategy should you take? A. Work with your ISP to diagnose the problem. B. Open a support ticket to ask for network capture and flow data to diagnose the problem, then roll back your application. C. Roll back to an earlier known good release initially, then use Stackdriver Trace and logging to diagnose the problem in a development/test/staging environment. D. Roll back to an earlier known good release, then push the release again at a quieter period to investigate. Then use Stackdriver Trace and logging to diagnose the problem. Answer: D NEW QUESTION 28 Your company has successfully migrated to the cloud and wants to analyze their data stream to optimize operations. They do not have any existing code for this analysis, so they are exploring all their options. These options include a mix of batch and stream processing, as they are running some hourly jobs and live-processing some data as it comes in. Which technology should they use for this? A. Google Cloud Dataproc B. Google Cloud Dataflow C. Google Container Engine with Bigtable D. Google Compute Engine with Google BigQuery Answer: B Explanation: Dataflow is for processing both the Batch and Stream. NEW QUESTION 29 Your customer is moving an existing corporate application to Google Cloud Platform from an on-premises data center. The business owners require minimal user disruption. There are strict security team requirements for storing passwords. What authentication strategy should they use?

A. Use G Suite Password Sync to replicate passwords into Google. B. Federate authentication via SAML 2.0 to the existing Identity Provider. C. Cloud Directory Sync tool. D. Ask users to set their Google password to match their corporate password. Answer: A Explanation: <u>https://support.google.com/a/answer/2611859?hl=en</u> NEW QUESTION 30 You are creating a solution to remove backup files older than 90 days from your backup Cloud Storage bucket. You want to optimize ongoing Cloud Storage spend. What should you do? A. Write a lifecycle management rule in XML and push it to the bucket with gsutil. B. Write a lifecycle management rule in JSON and push it to the bucket with gsutil. C. Schedule a cron script using gsutil is -lr gs://backups/** to find and remove items older than 90 days and schedule it with cron. Answer: B Explanation: <u>https://cloud.google.com/storage/docs/gsutil/commands/lifecycle</u> Download the newest PassLeader Professional Cloud Architect dumps from passleader.com now! 100% Pass Guarantee! Professional Cloud Architect VCE

dumps: https://www.passleader.com/professional-cloud-architect.html (82 Q&As) (New Questions Are 100% Available and Wrong Answers Have Been Corrected! Free VCE simulator!) P.S. New Professional Cloud Architect dumps PDF: https://drive.google.com/open?id=19jt3GbCmVz-pmGbZv8zjAu0NH7423IQ2 P.S. New Professional Data Engineer dumps PDF: https://drive.google.com/open?id=1m882ngsiRO1BOHineV4IQUv9jgF5Lpue