Amazon AWS-Solutions-Architect-Professional Upgrade Dumps, AWS-Solutions-Architect-Professional Exam Simulations

BONUS!!! Download part of Real4dumps AWS-Solutions-Architect-Professional dumps for free: https://drive.google.com/open?id=1lxOCqOOOm6R1S5S9EvwvhQbLzWhaSl7B

Real4dumps exam dumps have two version-PDF and SOFT version which will give you convenient. It is very convenient for you to use PDF real questions and answers. And you can download these materials and print it out for study at any time. The SOFT version simulates the real exam which will give you more realistic feeling. When you are faced with the real exam, you can pass Amazon AWS-Solutions-Architect-Professional test easily.

Amazon AWS-Solutions-Architect-Professional (AWS Certified Solutions Architect - Professional) Certification Exam is designed for professionals in the field of cloud computing who are seeking to enhance their skills and knowledge in the Amazon Web Services (AWS) platform. AWS Certified Solutions Architect - Professional certification exam is a validation of a candidate's expertise in designing and deploying scalable, highly available, and fault-tolerant systems on AWS.

>> Amazon AWS-Solutions-Architect-Professional Upgrade Dumps <<

AWS-Solutions-Architect-Professional Exam Simulations, Latest AWS-Solutions-Architect-Professional Test Pdf

The price for AWS-Solutions-Architect-Professional exam torrent is quite reasonable, you can afford it no matter you are a student or you are an employee in the company. You just need to spend some money, and you can get a certificate. In addition, AWS-Solutions-Architect-Professional exam dumps are high-quality and accuracy, and you can pass the exam successfully by using them. We also pass guarantee and money back guarantee for your failure of the exam after using AWS-Solutions-Architect-Professional Exam Dumps. We offer you free update for 365 days after purchasing, and the update version will be sent to your email address automatically.

To prepare for the AWS-Solutions-Architect-Professional Exam, candidates should have a solid understanding of AWS services and should be familiar with advanced topics such as security, networking, and database architecture. They should also be comfortable with designing and deploying complex systems in the cloud. Many candidates choose to enroll in training courses or study guides to help them prepare for the exam.

Amazon AWS Certified Solutions Architect - Professional Sample Questions (Q312-Q317):

NEW QUESTION # 312
A company ingests and processes streaming market data. The data rate is constant. A nightly process that calculates aggregate statistics is run, and each execution takes about 4 hours to complete. The statistical analysis is not mission critical to the business, and previous data points are picked up on the next execution if a particular run fails.
The current architecture uses a pool of Amazon EC2 Reserved Instances with 1-year reservations running full time to ingest and

store the streaming data in attached Amazon EBS volumes. On-Demand EC2 instances are launched each night to perform the nightly processing, accessing the stored data from NFS shares on the ingestion servers, and terminating the nightly processing servers when complete. The Reserved Instance reservations are expiring, and the company needs to determine whether to purchase new reservations or implement a new design.
Which is the most cost-effective design?

- A. Update the ingestion process to use Amazon Kinesis Data Firehose to save data to Amazon S3.
  Use a fleet of On-Demand EC2 instances that launches each night to perform the batch processing of the S3 data and terminates when the processing completes.
- B. Update the ingestion process to use Amazon Kinesis Data Firehouse to save data to Amazon S3.
  Use AWS Batch to perform nightly processing with a Spot market bid of 50% of the On-Demand price.
- C. Update the ingestion process to use a fleet of EC2 Reserved Instances behind a Network Load Balancer with 3-year leases. Use Batch with Spot instances with a maximum bid of 50% of the On- Demand price for the nightly processing.
- D. Update the ingestion process to use Amazon Kinesis Data Firehose to save data to Amazon Redshift.
  Use an AWS Lambda function scheduled to run nightly with Amazon CloudWatch Events to query Amazon Redshift to generate the daily statistics.

Answer: B

Explanation:
A: On demand instances is more expensive.
B: Since the analysis is not mission critical and can be restarted, spot instances is cheaper.
C: EC2 is more expensive.
D: Compared to B Lambda should be cheaper to Batch using spot instances which ultimately still charges based on EC2 pricing. But critical thing is Lambda cannot go beyond 15 minutes of execution time. So it is more for simple processing and in this case it is not. It takes 4 hours.
https://www.simform.com/aws-lambda-vs-ec2/
https://aws.amazon.com/about-aws/whats-new/2018/10/aws-lambda-supports-functions-that-can- run-up-to-15-minutes/
https://aws.amazon.com/batch/faqs/?nc=sn&loc=5

NEW QUESTION # 313
An on-premises application will be migrated to the cloud. The application consists of a single Elasticsearch virtual machine with data source feeds from local systems that will not be migrated, and a Java web application on Apache Tomcat running on three virtual machines. The Elasticsearch server currently uses 1 TB of storage out of 16 TB available storage, and the web application is updated every 4 months. Multiple users access the web application from the Internet. There is a 10Gbit AWS Direct Connect connection established, and the application can be migrated over a schedules 48-hour change window.
Which strategy will have the LEAST impact on the Operations staff after the migration?

- A. Create an Amazon ES cluster for Elasticsearch and a public AWS Elastic Beanstalk environment for the web application. Pause the data source feeds, export the Elasticsearch index from on premises, and import into the Amazon ES cluster. Move the data source feeds to the new Amazon ES cluster endpoint and move users to the new web application.
- B. Create an Amazon ES cluster for Elasticsearch and a public AWS Elastic Beanstalk environment for the web application. Use AWS DMS to replicate Elasticsearch data. When replication has finished, move data source feeds to the new Amazon ES cluster endpoint and move users to the new web application.
- C. Create an Elasticsearch server on Amazon EC2 right-sized with 2 TB of Amazon EBS and a public AWS Elastic Beanstalk environment for the web application. Pause the data sources, export the Elasticsearch index from on premises, and import into the EC2 Elasticsearch server.
  Move data source feeds to the new Elasticsearch server and move users to the web application.
- D. Use the AWS SMS to replicate the virtual machines into AWS. When the migration is complete, pause the data source feeds and start the migrated Elasticsearch and web application instances.
  Place the web application instances behind a public Elastic Load Balancer. Move the data source feeds to the new Elasticsearch server and move users to the new web Application Load Balancer.

Answer: B

Explanation:
A: uses EC2s, and that requires maintenance after the migration. contradicting with least operation after migration.
C: Good, but again uses self Elastic Search and EC2s, which requires operation maintenance staff after migration.
D: the sequence of moving the index to ES is wrong. They should have replicate then pause!
However, answer "B" uses DMS as a replication tool.

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.Elasticsearch.html
https://aws.amazon.com/blogs/database/introducing-amazon-elasticsearch-service-as-a-target-in- aws-database-migration-service/

NEW QUESTION # 314
A company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. On the first day of every month at midnight, the application becomes much slower when the month-end financial calculation batch executes. This causes the CPU utilization of the EC2 instances to immediately peak to 100%, which disrupts the application.
What should a solutions architect recommend to ensure the application is able to handle the workload and avoid downtime?

- A. Configure an EC2 Auto Scaling scheduled scaling policy based on the monthly schedule.
- B. Configure an EC2 Auto Scaling simple scaling policy based on CPU utilization.
- C. Configure an Amazon CloudFront distribution in front of the ALB.
- D. Configure Amazon ElastiGache to remove some of the workload from the EC2 instances.

Answer: A

NEW QUESTION # 315
A company has an on-premises volume backup solution that has reached its end of life. The company wants to use AWS as part of a new backup solution and wants to maintain local access to all the data while it is backed up on AWS. The company wants to ensure that the data backed up on AWS is automatically and securely transferred.
Which solution meets these requirements?

- A. Use AWS Snowball to migrate data out of the on-premises solution to Amazon S3. Configure on- premises systems to mount the Snowball S3 endpoint to provide local access to the data.
- B. Use AWS Storage Gateway and configure a stored volume gateway. Run the Storage Gateway software appliance on premises and map the gateway storage volumes to on-premises storage. Mount the gateway storage volumes to provide local access to the data.
- C. Use AWS Snowball Edge to migrate data out of the on-premises solution to Amazon S3. Use the Snowball Edge file interface to provide on-premises systems with local access to the data.
- D. Use AWS Storage Gateway and configure a cached volume gateway. Run the Storage Gateway software appliance on premises and configure a percentage of data to cache locally. Mount the gateway storage volumes to provide local access to the data.

Answer: B

NEW QUESTION # 316
An organization is making software for the CIA in USA. CIA agreed to host the application on AWS but in a secure environment. The organization is thinking of hosting the application on the AWS GovCloud region.
Which of the below mentioned difference is not correct when the organization is hosting on the AWS GovCloud in comparison with the AWS standard region?

- A. Physical and logical administrative access only to U.S. persons.
- B. The billing for the AWS GovCLoud will be in a different account than the Standard AWS account.
- C. It is physically isolated and has logical network isolation from all the other regions.
- D. GovCloud region authentication is isolated from Amazon.com.

Answer: B

Explanation:
Explanation
AWS GovCloud (US) is an isolated AWS region designed to allow U.S. government agencies and customers to move sensitive workloads into the cloud by addressing their specific regulatory and compliance requirements. The AWS GovCloud (US) Region adheres to the U.S. International Traffic in Arms Regulations (ITAR) requirements. It has added advantages, such as:
Restricting physical and logical administrative access to U.S. persons only There will be a separate AWS GovCloud (US) credentials, such as access key and secret access key than the standard AWS account The user signs in with the IAM user name and password The AWS GovCloud (US) Region authentication is completely isolated from Amazon.com If the organization is

planning to host on EC2 in AWS GovCloud then it will be billed to standard AWS account of organization since AWS GovCloud billing is linked with the standard AWS account and is not be billed separately.
http://docs.aws.amazon.com/govcloud-us/latest/UserGuide/whatis.html

NEW QUESTION # 317
.....