

Exam Questions SAP-C01

AWS Certified Solutions Architect- Professional

<https://www.2passeasy.com/dumps/SAP-C01/>



NEW QUESTION 1

A Solutions Architect must design a highly available, stateless, REST service. The service will require multiple persistent storage layers for service object meta information and the delivery of content. Each request needs to be authenticated and securely processed. There is a requirement to keep costs as low as possible? How can these requirements be met?

- A. Use AWS Fargate to host a container that runs a self-contained REST service
- B. Set up an Amazon ECS service that is fronted by an Application Load Balancer (ALB). Use a custom authenticator to control access to the AP
- C. Store request meta information in Amazon DynamoDB with Auto Scaling and static content in a secured S3 bucket
- D. Make secure signed requests for Amazon S3 objects and proxy the data through the REST service interface.
- E. Use AWS Fargate to host a container that runs a self-contained REST service
- F. Set up an ECS service that is fronted by a cross-zone AL
- G. Use an Amazon Cognito user pool to control access to the AP
- H. Store request meta information in DynamoDB with Auto Scaling and static content in a secured S3 bucket
- I. Generate presigned URLs when returning references to content stored in Amazon S3.
- J. Set up Amazon API Gateway and create the required API resources and method
- K. Use an Amazon Cognito user pool to control access to the AP
- L. Configure the methods to use AWS Lambda proxy integrations, and process each resource with a unique AWS Lambda function
- M. Store request meta information in DynamoDB with Auto Scaling and static content in a secured S3 bucket
- N. Generate presigned URLs when returning references to content stored in Amazon S3.
- O. Set up Amazon API Gateway and create the required API resources and method
- P. Use an Amazon API Gateway custom authorizer to control access to the AP
- Q. Configure the methods to use AWS Lambda custom integrations, and process each resource with a unique Lambda function
- R. Store request meta information in an Amazon ElastiCache Multi-AZ cluster and static content in a secured S3 bucket
- S. Generate presigned URLs when returning references to content stored in Amazon S3.

Answer: C

NEW QUESTION 2

A large company has many business units. Each business unit has multiple AWS accounts for different purposes. The CIO of the company sees that each business unit has data that would be useful to share with other parts of the company. In total, there are about 10 PB of data that needs to be shared with users in 1,000 AWS accounts. The data is proprietary, so some of it should only be available to users with specific job types. Some of the data is used for throughput of intensive workloads, such as simulations. The number of accounts changes frequently because of new initiatives, acquisitions, and divestitures. A Solutions Architect has been asked to design a system that will allow for sharing data for use in AWS with all of the employees in the company. Which approach will allow for secure data sharing in a scalable way?

- A. Store the data in a single Amazon S3 bucket
- B. Create an IAM role for every combination of job type and business unit that allows to appropriate read/write access based on object prefixes in the S3 bucket
- C. The roles should have trust policies that allow the business unit's AWS accounts to assume their role
- D. Use IAM in each business unit's AWS account to prevent them from assuming roles for a different job type
- E. Users get credentials to access the data by using AssumeRole from their business unit's AWS account
- F. Users can then use those credentials with an S3 client.
- G. Store the data in a single Amazon S3 bucket
- H. Write a bucket policy that uses conditions to grant read and write access where appropriate, based on each user's business unit and job type
- I. Determine the business unit with the AWS account accessing the bucket and the job type with a prefix in the IAM user's name
- J. Users can access data by using IAM credentials from their business unit's AWS account with an S3 client.
- K. Store the data in a series of Amazon S3 buckets
- L. Create an application running in Amazon EC2 that is integrated with the company's identity provider (IdP) that authenticates users and allows them to download or upload data through the application
- M. The application uses the business unit and job type information in the IdP to control what users can upload and download through the application
- N. The users can access the data through the application's API.
- O. Store the data in a series of Amazon S3 buckets
- P. Create an AWS STS token vending machine that is integrated with the company's identity provider (IdP). When a user logs in, have the token vending machine attach an IAM policy that assumes the role that limits the user's access and/or upload only the data the user is authorized to access
- Q. Users can get credentials by authenticating to the token vending machine's website or API and then use those credentials with an S3 client.

Answer: B

NEW QUESTION 3

A company currently uses Amazon EBS and Amazon RDS for storage purposes. The company intends to use a pilot light approach for disaster recovery in a different AWS Region. The company has an RTO of 6 hours and an RPO of 24 hours. Which solution would achieve the requirements with MINIMAL cost?

- A. Use AWS Lambda to create daily EBS and RDS snapshots, and copy them to the disaster recovery region
- B. Use Amazon Route 53 with active-passive failover configuration
- C. Use Amazon EC2 in an Auto Scaling group with the capacity set to 0 in the disaster recovery region.
- D. Use AWS Lambda to create daily EBS and RDS snapshots, and copy them to the disaster recovery region
- E. Use Amazon Route 53 with active-active failover configuration
- F. Use Amazon EC2 in an AutoScaling group configured in the same way as in the primary region.
- G. Use Amazon ECS to handle long-running tasks to create daily EBS and RDS snapshots, and copy to the disaster recovery region
- H. Use Amazon Route 53 with active-passive failover configuration
- I. Use Amazon EC2 in an Auto Scaling group with the capacity set to 0 in the disaster recovery region.
- J. Use EBS and RDS cross-region snapshot copy capability to create snapshots in the disaster recovery region
- K. Use Amazon Route 53 with active-active failover configuration
- L. Use Amazon EC2 in an Auto Scaling group with the capacity set to 0 in the disaster recovery region.

Answer: A

Explanation:

https://docs.aws.amazon.com/AmazonECS/latest/developerguide/scheduling_tasks.html

NEW QUESTION 4

A company had a tight deadline to migrate its on-premises environment to AWS. It moved over Microsoft SQL Servers and Microsoft Windows Servers using the virtual machine import/export service and rebuild other applications native to the cloud. The team created both Amazon EC2 databases and used Amazon RDS. Each team in the company was responsible for migrating their applications, and they have created individual accounts for isolation of resources. The company did not have much time to consider costs, but now it would like suggestions on reducing its AWS spend.

Which steps should a Solutions Architect take to reduce costs?

- A. Enable AWS Business Support and review AWS Trusted Advisor's cost check
- B. Create Amazon EC2 Auto Scaling groups for applications that experience fluctuating demand
- C. Save AWS Simple Monthly Calculator reports in Amazon S3 for trend analysis
- D. Create a master account under Organizations and have teams join for consolidating billing.
- E. Enable Cost Explorer and AWS Business Support Reserve Amazon EC2 and Amazon RDS DB instance
- F. Use Amazon CloudWatch and AWS Trusted Advisor for monitoring and to receive cost-savings suggestion
- G. Create a master account under Organizations and have teams join for consolidated billing.
- H. Create an AWS Lambda function that changes the instance size based on Amazon CloudWatch alarms. Reserve instances based on AWS Simple Monthly Calculator suggestion
- I. Have an AWS Well-Architected framework review and apply recommendation
- J. Create a master account under Organizations and have teams join for consolidated billing.
- K. Create a budget and monitor for costs exceeding the budget
- L. Create Amazon EC2 Auto Scaling groups for applications that experience fluctuating demand
- M. Create an AWS Lambda function that changes instance sizes based on Amazon CloudWatch alarm
- N. Have each team upload their bill to an Amazon S3 bucket for analysis of team spending
- O. Use Spot instances on nightly batch processing jobs.

Answer: D

NEW QUESTION 5

While debugging a backend application for an IoT system that supports globally distributed devices a Solutions Architect notices that stale data is occasionally being sent to user devices. Devices often share data, and stale data does not cause issues in most cases. However, device operations are disrupted when a device reads the stale data after an update.

The global system has multiple identical application stacks deployed in different AWS Regions. If a user device travels out of its home geographic region, it will always connect to the geographically closest AWS Region to write or read data. The same data is available in all supported AWS Regions using an Amazon DynamoDB global table.

What change should be made to avoid causing disruptions in device operations?

- A. Update the backend to use strongly consistent read
- B. Update the devices to always write to and read from their home AWS Region
- C. Enable strong consistency globally on a DynamoDB global table. Update the backend to use strongly consistent reads
- D. Switch the backend data store to Amazon Aurora MySQL with cross-region replicas. Update the backend to always write to the master endpoint
- E. Select one AWS Region as a master and perform all writes in that AWS Region only. Update the backend to use strongly consistent reads

Answer: B

NEW QUESTION 6

A company has a requirement that only allows specially hardened AMIs to be launched into public subnets in a VPC, and for the AMIs to be associated with a specific security group. Allowing non-compliant instances to launch into the public subnet could present a significant security risk if they are allowed to operate. A mapping of approved AMIs to subnets to security groups exists in an Amazon DynamoDB table in the same AWS account. The company created an AWS Lambda function that, when invoked, will terminate a given Amazon EC2 instance if the combination of AMI, subnet, and security group are not approved in the DynamoDB table.

What should the Solutions Architect do to MOST quickly mitigate the risk of compliance deviations?

- A. Create an Amazon CloudWatch Events rule that matches each time an EC2 instance is launched using one of the allowed AMIs, and associate it with the Lambda function as the target.
- B. For the Amazon S3 bucket receiving the AWS CloudTrail logs, create an S3 event notification configuration with a filter to match when logs contain the `ec2:RunInstances` action, and associate it with the Lambda function as the target.
- C. Enable AWS CloudTrail and configure it to stream to an Amazon CloudWatch Logs group
- D. Create a metric filter in CloudWatch to match when the `ec2:RunInstances` action occurs, and trigger the Lambda function when the metric is greater than 0.
- E. Create an Amazon CloudWatch Events rule that matches each time an EC2 instance is launched, and associate it with the Lambda function as the target.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-lifecycle.html>

NEW QUESTION 7

A company is migrating its marketing website and content management system from an on-premises data center to AWS. The company wants the AWS application to be developed in a VPC with Amazon EC2 instances used for the web servers and an Amazon RDS instance for the database.

The company has a runbook document that describes the installation process of the on-premises system. The company would like to base the AWS system on the processes referenced in the runbook document. The runbook document describes the installation and configuration of the operating systems, network settings, the website, and content management system software on the servers. After the migration is complete, the company wants to be able to make changes quickly to take advantage of other AWS features.

How can the application and environment be deployed and automated in AWS, while allowing for future changes?

- A. Update the runbook to describe how to create the VPC, the EC2 instances, and the RDS instance for the application by using the AWS Console
- B. Make sure that the rest of the steps in the runbook are updated to reflect any changes that may come from the AWS migration.
- C. Write a Python script that uses the AWS API to create the VPC, the EC2 instances, and the RDS instance for the application
- D. Write shell scripts that implement the rest of the steps in the runbook
- E. Have the Python script copy and run the shell scripts on the newly created instances to complete the installation.
- F. Write an AWS CloudFormation template that creates the VPC, the EC2 instances, and the RDS instance for the application

- G. Ensure that the rest of the steps in the runbook are updated to reflect any changes that may come from the AWS migration.
- H. Write an AWS CloudFormation template that creates the VPC, the EC2 instances, and the RDS instance for the applicatio
- I. Include EC2 user data in the AWS CloudFormation template to install and configure the software.

Answer: D

NEW QUESTION 8

A company is using AWS CloudFormation to deploy its infrastructure. The company is concerned that, if a production CloudFormation stack is deleted, important data stored in Amazon RDS databases or Amazon EBS volumes might also be deleted. How can the company prevent users from accidentally deleting data in this way?

- A. Modify the CloudFormation templates to add a DeletionPolicy attribute to RDS and EBS resources.
- B. Configure a stack policy that disallows the deletion of RDS and EBS resources.
- C. Modify IAM policies to deny deleting RDS and EBS resources that are tagged with an "aws:cloudformation:stack-name" tag.
- D. Use AWS Config rules to prevent deleting RDS and EBS resources.

Answer: A

Explanation:

With the DeletionPolicy attribute you can preserve or (in some cases) backup a resource when its stack is deleted. You specify a DeletionPolicy attribute for each resource that you want to control. If a resource has no DeletionPolicy attribute, AWS CloudFormation deletes the resource by default. To keep a resource when its stack is deleted, specify Retain for that resource. You can use retain for any resource. For example, you can retain a nested stack, Amazon S3 bucket, or EC2 instance so that you can continue to use or modify those resources after you delete their stacks.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html>

NEW QUESTION 9

A company will several AWS accounts is using AWS Organizations and service control policies (SCPs). An Administrator created the following SCP and has attached it to an organizational unit (OU) that contains AWS account 1111-1111-1111:

```
{
  "Version": "2012-10-27",
  "Statement": [
    {
      "Sid": "AllowsAllActions",
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*"
    },
    {
      "Sid": "DenyCloudTrail",
      "Effect": "Deny",
      "Action": "cloudtrail:*",
      "Resource": "*"
    }
  ]
}
```

Developers working in account 1111-1111-1111 complain that they cannot create Amazon S3 buckets. How should the Administrator address this problem?

- A. Add s3:CreateBucket with "Allow" effect to the SCP.
- B. Remove the account from the OU, and attach the SCP directly to account 1111-1111-1111.
- C. Instruct the Developers to add Amazon S3 permissions to their IAM entities.
- D. Remove the SCP from account 1111-1111-1111.

Answer: C

NEW QUESTION 10

A company runs a three-tier application in AWS. Users report that the application performance can vary greatly depending on the time of day and functionality being accessed.

The application includes the following components:

- Eight t2.large front-end web servers that serve static content and proxy dynamic content from the application tier.
- Four t2.large application servers.
- One db.m4.large Amazon RDS MySQL Multi-AZ DB instance.

Operations has determined that the web and application tiers are network constrained.

Which of the following should cost effective improve application performance? (Choose two.)

- A. Replace web and app tiers with t2.xlarge instances
- B. Use AWS Auto Scaling and m4.large instances for the web and application tiers
- C. Convert the MySQL RDS instance to a self-managed MySQL cluster on Amazon EC2
- D. Create an Amazon CloudFront distribution to cache content
- E. Increase the size of the Amazon RDS instance to db.m4.xlarge

Answer: BD

Explanation:

<https://aws.amazon.com/ec2/instance-types/>

NEW QUESTION 10

A large multinational company runs a timesheet application on AWS that is used by staff across the world. The application runs on Amazon EC2 instances in an Auto Scaling group behind an Elastic Load Balancing (ELB) load balancer, and stores in an Amazon RDS MySQL Multi-AZ database instance. The CFO is concerned about the impact on the business if the application is not available. The application must not be down for more than two hours, but the solution must be as cost-effective as possible.

How should the Solutions Architect meet the CFO's requirements while minimizing data loss?

- A. In another region, configure a read replica and create a copy of the infrastructure
- B. When an issue occurs, promote the read replica and configure as an Amazon RDS Multi-AZ database instance
- C. Update the DNS to point to the other region's ELB.
- D. Configure a 1-day window of 60-minute snapshots of the Amazon RDS Multi-AZ database instance. Create an AWS CloudFormation template of the application infrastructure that uses the latest snapshot
- E. When an issue occurs, use the AWS CloudFormation template to create the environment in another region
- F. Update the DNS record to point to the other region's ELB.
- G. Configure a 1-day window of 60-minute snapshots of the Amazon RDS Multi-AZ database instance which is copied to another region
- H. Create an AWS CloudFormation template of the application infrastructure that uses the latest copied snapshot
- I. When an issue occurs, use the AWS CloudFormation template to create the environment in another region
- J. Update the DNS record to point to the other region's ELB.
- K. Configure a read replica in another region
- L. Create an AWS CloudFormation template of the application infrastructure
- M. When an issue occurs, promote the read replica and configure as an Amazon RDS Multi-AZ database instance and use the AWS CloudFormation template to create the environment in another region using the promoted Amazon RDS instance
- N. Update the DNS record to point to the other region's ELB.

Answer: D

NEW QUESTION 11

A company is migrating an application to AWS. It wants to use fully managed services as much as possible during the migration. The company needs to store large, important documents within the application with the following requirements:

- The data must be highly durable and available.
- The data must always be encrypted at rest and in transit.
- The encryption key must be managed by the company and rotated periodically. Which of the following solutions should the Solutions Architect recommend?

- A. Deploy the storage gateway to AWS in file gateway mode
- B. Use Amazon EBS volume encryption using an AWS KMS key to encrypt the storage gateway volumes.
- C. Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.
- D. Use Amazon DynamoDB with SSL to connect to DynamoDB
- E. Use an AWS KMS key to encrypt DynamoDB objects at rest.
- F. Deploy instances with Amazon EBS volumes attached to store this data
- G. Use EBS volume encryption using an AWS KMS key to encrypt the data.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/security/how-to-use-bucket-policies-and-apply-defense-in-depth-to-help-secure-y>

NEW QUESTION 14

A company's CISO has asked a Solutions Architect to re-engineer the company's current CI/CD practices to make sure patch deployments to its applications can happen as quickly as possible with minimal downtime if vulnerabilities are discovered. The company must also be able to quickly roll back a change in case of errors. The web application is deployed in a fleet of Amazon EC2 instances behind an Application Load Balancer. The company is currently using GitHub to host the application source code and has configured an AWS CodeBuild project to build the application. The company also intends to use AWS CodePipeline to trigger builds from GitHub commits using the existing CodeBuild project.

What CI/CD configuration meets all of the requirements?

- A. Configure CodePipeline with a deploy stage using AWS CodeDeploy configured for in-place deployment
- B. Monitor the newly deployed code, and if there are any issues, push another code update.
- C. Configure CodePipeline with a deploy stage using AWS CodeDeploy configured for blue/green deployment
- D. Monitor the new deployed code and if there are any issues, trigger a manual rollback using CodeDeploy.
- E. Configure CodePipeline with a deploy stage using AWS CloudFormation to create a pipeline for test and production stack
- F. Monitor the newly deployed code and if there are any issues push another code update.
- G. Configure the CodePipeline with a deploy stage using AWS OpsWorks and in-place deployments. Monitor the newly deployed code and if there are any issues, push another code update.

Answer: B

NEW QUESTION 19

An online e-commerce business is running a workload on AWS. The application architecture includes a web tier, an application tier for business logic, and a database tier for user and transactional data management. The database server has a 100 GB memory requirement. The business requires cost-efficient disaster recovery for the application with an RTO of 5 minutes and an RPO of 1 hour. The business also has a regulatory requirement for out-of-region disaster recovery with a minimum distance between the primary and alternate sites of 250 miles.

Which of the following options can the Solutions Architect design to create a comprehensive solution for this customer that meets the disaster recovery requirements?

- A. Back up the application and database data frequently and copy them to Amazon S3. Replicate the backups using S3 cross-region replication, and use AWS CloudFormation to instantiate infrastructure for disaster recovery and restore data from Amazon S3.
- B. Employ a pilot light environment in which the primary database is configured with mirroring to build a standby database on m4.large in the alternate regio
- C. Use AWS CloudFormation to instantiate the web servers, application servers and load balancers in case of a disaster to bring the application up in the alternate regio
- D. Vertically resize the database to meet the full production demands, and use Amazon Route 53 to switch traffic to the alternate region.
- E. Use a scaled-down version of the fully functional production environment in the alternate region that includes one instance of the web server, one instance of the application server, and a replicated instance of the database server in standby mod
- F. Place the web and the application tiers in an Auto Scaling behind a load balancer, which can automatically scale when the load arrives to the applicatio
- G. Use Amazon Route 53 to switch traffic to the alternate region.
- H. Employ a multi-region solution with fully functional web, application, and database tiers in both regions with equivalent capacit
- I. Activate the primary database in one region only and the standby database in the other regio
- J. Use Amazon Route 53 to automatically switch traffic from one region to another using health check routing policies.

Answer: C

NEW QUESTION 23

A company deployed a three-tier web application in two regions: us-east-1 and eu-west-1. The application must be active in both regions at the same time. The database tier of the application uses a single Amazon RDS Aurora database globally, with a master in us-east-1 and a read replica in eu-west-1. Both regions are connected by a VPN.

The company wants to ensure that the application remains available even in the event of a region-level failure of all of the application's components. It is acceptable for the application to be in read-only mode for up to 1 hour. The company plans to configure two Amazon Route 53 record sets, one for each of the regions.

How should the company complete the configuration to meet its requirements while providing the lowest latency for the application end-users? (Choose two.)

- A. Use failover routing and configure the us-east-1 record set as primary and the eu-west-1 record set as secondar
- B. Configure an HTTP health check for the web application in us-east-1, and associate it to the us-east-1 record set.
- C. Use weighted routing and configure each record set with a weight of 50. Configure an HTTP health check for each region, and attach it to the record set for that region.
- D. Use latency-based routing for both record set
- E. Configure a health check for each region and attach it to the record set for that region.
- F. Configure an Amazon CloudWatch alarm for the health checks in us-east-1, and have it invoke an AWS Lambda function that promotes the read replica in eu-west-1.
- G. Configure an Amazon RDS event notifications to react to the failure of the database in us-east-1 by invoking an AWS Lambda function that promotes the read replica in eu-west-1.

Answer: CE

Explanation:

<https://docs.aws.amazon.com/lambda/latest/dg/services-rds.html>

NEW QUESTION 25

A company runs a legacy system on a single m4.2xlarge Amazon EC2 instance with Amazon EBS2 storage. The EC2 instance runs both the web server and a self-managed Oracle database. A snapshot is made of the EBS volume every 12 hours, and an AMI was created from the fully configured EC2 instance.

A recent event that terminated the EC2 instance led to several hours of downtime. The application was successfully launched from the AMI, but the age of the EBS snapshot and the repair of the database resulted in the loss of 8 hours of data. The system was also down for 4 hours while the Systems Operators manually performed these processes.

What architectural changes will minimize downtime and reduce the chance of lost data?

- A. Create an Amazon CloudWatch alarm to automatically recover the instanc
- B. Create a script that will check and repair the database upon reboo
- C. Subscribe the Operations team to the Amazon SNS message generated by the CloudWatch alarm.
- D. Run the application on m4.xlarge EC2 instances behind an Elastic Load Balancer/Application Load Balance
- E. Run the EC2 instances in an Auto Scaling group across multiple Availability Zones with a minimum instance count of tw
- F. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.
- G. Run the application on m4.2xlarge EC2 instances behind an Elastic Load Balancer/Application Load Balance
- H. Run the EC2 instances in an Auto Scaling group across multiple Availability Zones with aminimum instance count of on
- I. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.
- J. Increase the web server instance count to two m4.xlarge instances and use Amazon Route 53 round-robin load balancing to spread the loa
- K. Enable Route 53 health checks on the web server
- L. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.

Answer: B

Explanation:

Ensures that there are at least two EC instances, each of which is in a different AZ. It also ensures that the database spans multiple AZs. Hence this meets all the criteria.

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

NEW QUESTION 30

A Solutions Architect has been asked to look at a company's Amazon Redshift cluster, which has quickly become an integral part of its technology and supports key business process. The Solutions Architect is to increase the reliability and availability of the cluster and provide options to ensure that if an issue arises, the cluster can either operate or be restored within four hours.

Which of the following solution options BEST addresses the business need in the most cost-effective manner?

- A. Ensure that the Amazon Redshift cluster has been set up to make use of Auto Scaling groups with the nodes in the cluster spread across multiple Availability Zones.
- B. Ensure that the Amazon Redshift cluster creation has been template using AWS CloudFormation so it can easily be launched in another Availability Zone and data populated from the automated Redshift back-ups stored in Amazon S3.
- C. Use Amazon Kinesis Data Firehose to collect the data ahead of ingestion into Amazon Redshift and create clusters using AWS CloudFormation in another

region and stream the data to both clusters.

D. Create two identical Amazon Redshift clusters in different regions (one as the primary, one as the secondary). Use Amazon S3 cross-region replication from the primary to secondary). Use Amazon S3 cross-region replication from the primary to secondary region, which triggers an AWS Lambda function to populate the cluster in the secondary region.

Answer: B

Explanation:

https://aws.amazon.com/redshift/faqs/?nc1=h_ls Q: What happens to my data warehouse cluster availability and data durability if my data warehouse cluster's Availability Zone (AZ) has an outage? If your Amazon Redshift data warehouse cluster's Availability Zone becomes unavailable, you will not be able to use your cluster until power and network access to the AZ are restored. Your data warehouse cluster's data is preserved so you can start using your Amazon Redshift data warehouse as soon as the AZ becomes available again. In addition, you can also choose to restore any existing snapshots to a new AZ in the same Region. Amazon Redshift will restore your most frequently accessed data first so you can resume queries as quickly as possible.

FROM 37

NEW QUESTION 34

A company is moving a business-critical application onto AWS. It is a traditional three-tier web application using an Oracle database. Data must be encrypted in transit and at rest. The database hosts 12 TB of data. Network connectivity to the source Oracle database over the internal is allowed, and the company wants to reduce the operational costs by using AWS Managed Services where possible. All resources within the web and application tiers have been migrated. The database has a few tables and a simple schema using primary keys only; however, it contains many Binary Large Object (BLOB) fields. It was not possible to use the database's native replication tools because of licensing restrictions.

Which database migration solution will result in the LEAST amount of impact to the application's availability?

- A. Provision an Amazon RDS for Oracle instance
- B. Host the RDS database within a virtual private cloud (VPC) subnet with internet access, and set up the RDS database as an encrypted Read Replica of the source databases
- C. Use SSL to encrypt the connection between the two databases
- D. Monitor the replication performance by watching the RDS ReplicaLag metric
- E. During the application maintenance window, shut down the on-premises database and switch over the application connection to the RDS instance when there is no more replication lag
- F. Promote the Read Replica into a standalone database instance.
- G. Provision an Amazon EC2 instance and install the same Oracle database software
- H. Create a backup of the source database using the supported tool
- I. During the application maintenance window, restore the backup into the Oracle database running in the EC2 instance
- J. Set up an Amazon RDS for Oracle instance, and create an import job between the database hosted in AWS
- K. Shut down the source database and switch over the database connections to the RDS instance when the job is complete.
- L. Use AWS DMS to load and replicate the dataset between the on-premises Oracle database and the replication instance hosted on AWS
- M. Provision an Amazon RDS for Oracle instance with Transparent Data Encryption (TDE) enabled and configure it as target for the replication instance
- N. Create a customer-managed AWS KMS master key to set it as the encryption key for the replication instance. Use AWS DMS tasks to load the data into the target RDS instance
- O. During the application maintenance window and after the load tasks reach the ongoing replication phase, switch the database connections to the new database.
- P. Create a compressed full database backup on the on-premises Oracle database during an application maintenance window
- Q. While the backup is being performed, provision a 10 Gbps AWS Direct Connect connection to increase the transfer speed of the database backup files to Amazon S3, and shorten the maintenance window period
- R. Use SSL/TLS to copy the files over the Direct Connect connection
- S. When the backup files are successfully copied, start the maintenance window, and use any of the Amazon RDS supported tools to import the data into a newly provisioned Amazon RDS for Oracle instance with encryption enabled
- T. Wait until the data is fully loaded and switch over the database connections to the new databases
- . Delete the Direct Connect connection to cut unnecessary charges.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/apn/oracle-database-encryption-options-on-amazon-rds/>

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Appendix.Oracle.Options.AdvSecurity.htm> | (DMS in transit encryption)

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Security.html

NEW QUESTION 35

A company operating a website on AWS requires high levels of scalability, availability and performance. The company is running a Ruby on Rails application on Amazon EC2. It has a data tier on MySQL 5.6 on Amazon EC2 using 16 TB of Amazon EBS storage. Amazon CloudFront is used to cache application content. The Operations team is reporting continuous and unexpected growth of EBS volumes assigned to the MySQL database. The Solutions Architect has been asked to design a highly scalable, highly available, and high-performing solution.

Which solution is the MOST cost-effective at scale?

- A. Implement Multi-AZ and Auto Scaling for all EC2 instances in the current configuration
- B. Ensure that all EC2 instances are purchased as reserved instances
- C. Implement new elastic Amazon EBS volumes for the data tier.
- D. Design and implement the Docker-based containerized solution for the application using Amazon EC2
- E. Migrate to an Amazon Aurora MySQL Multi-AZ cluster
- F. Implement storage checks for Aurora MySQL storage utilization and an AWS Lambda function to grow the Aurora MySQL storage, as necessary
- G. Ensure that Multi-AZ architectures are implemented.
- H. Ensure that EC2 instances are right-sized and behind an Elastic Load Balancing load balancer. Implement Auto Scaling with EC2 instances
- I. Ensure that the reserved instances are purchased for fixed capacity and that Auto Scaling instances run on demand
- J. Migrate to an Amazon Aurora MySQL Multi-AZ cluster
- K. Ensure that Multi-AZ architectures are implemented.
- L. Ensure that EC2 instances are right-sized and behind an Elastic Load Balance
- M. Implement Auto Scaling with EC2 instances
- N. Ensure that Reserved instances are purchased for fixed capacity and that Auto Scaling instances run on demand
- O. Migrate to an Amazon Aurora MySQL Multi-AZ cluster
- P. Implement storage checks for Aurora MySQL storage utilization and an AWS Lambda function to grow Aurora MySQL storage, as necessary
- Q. Ensure Multi-AZ architectures are implemented.

Answer: C

NEW QUESTION 37

An e-commerce company is revamping its IT infrastructure and is planning to use AWS services. The company's CIO has asked a Solutions Architect to design a simple, highly available, and loosely coupled order processing application. The application is responsible for receiving and processing orders before storing them in an Amazon DynamoDB table. The application has a sporadic traffic pattern and should be able to scale during marketing campaigns to process the orders with minimal delays.

Which of the following is the MOST reliable approach to meet the requirements?

- A. Receive the orders in an Amazon EC2-hosted database and use EC2 instances to process them.
- B. Receive the orders in an Amazon SQS queue and trigger an AWS Lambda function to process them.
- C. Receive the orders using the AWS Step Functions program and trigger an Amazon ECS container to process them.
- D. Receive the orders in Amazon Kinesis Data Streams and use Amazon EC2 instances to process them.

Answer: B

NEW QUESTION 41

A company has a large on-premises Apache Hadoop cluster with a 20 PB HDFS database. The cluster is growing every quarter by roughly 200 instances and 1 PB. The company's goals are to enable resiliency for its Hadoop data, limit the impact of losing cluster nodes, and significantly reduce costs. The current cluster runs 24/7 and supports a variety of analysis workloads, including interactive queries and batch processing.

Which solution would meet these requirements with the LEAST expense and down time?

- A. Use AWS Snowmobile to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workload based on historical data from the on-premises cluster.
- B. Store the data on EMRFS.
- C. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metric.
- D. Create job-specific, optimized clusters for batch workloads that are similarly optimized.
- E. Use AWS Snowmobile to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster of similar size and configuration to the current cluster.
- F. Store the data on EMRFS.
- G. Minimize costs by using Reserved Instance.
- H. As the workload grows each quarter, purchase additional Reserved Instances and add to the cluster.
- I. Use AWS Snowball to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workloads based on historical data from the on-premises cluster.
- J. Store the data on EMRFS.
- K. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metric.
- L. Create job-specific, optimized clusters for batch workloads that are similarly optimized.
- M. Use AWS Direct Connect to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workload based on historical data from the on-premises cluster.
- N. Store the data on EMRFS.
- O. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metric.
- P. Create job-specific, optimized clusters for batch workloads that are similarly optimized.

Answer: A

Explanation:

Q: How should I choose between Snowmobile and Snowball?

To migrate large datasets of 10PB or more in a single location, you should use Snowmobile. For datasets less than 10PB or distributed in multiple locations, you should use Snowball. In addition, you should evaluate the amount of available bandwidth in your network backbone. If you have a high speed backbone with hundreds of Gb/s of spare throughput, then you can use Snowmobile to migrate the large datasets all at once. If you have limited bandwidth on your backbone, you should consider using multiple Snowballs to migrate the data incrementally.

NEW QUESTION 46

A company wants to replace its call system with a solution built using AWS managed services. The company call center would like the solution to receive calls, create contact flows, and scale to handle growth projections. The call center would also like the solution to use deep learning capabilities to recognize the intent of the callers and handle basic tasks, reducing the need to speak an agent. The solution should also be able to query business applications and provide relevant information back to calls as requested.

Which services should the Solution Architect use to build this solution? (Choose three.)

- A. Amazon Rekognition to identify who is calling.
- B. Amazon Connect to create a cloud-based contact center.
- C. Amazon Alexa for Business to build conversational interface.
- D. AWS Lambda to integrate with internal systems.
- E. Amazon Lex to recognize the intent of the caller.
- F. Amazon SQS to add incoming callers to a queue.

Answer: BDE

NEW QUESTION 47

A company's main intranet page has experienced degraded response times as its user base has increased although there are no reports of users seeing error pages. The application uses Amazon DynamoDB in read-only mode.

Amazon DynamoDB latency metrics for successful requests have been in a steady state even during times when users have reported degradation. The Development team has correlated the issue to ProvisionedThroughputExceeded exceptions in the application logs when doing Scan and read operations. The team also identified an access pattern of steady spikes of read activity on a distributed set of individual data items.

The Chief Technology Officer wants to improve the user experience.

Which solutions will meet these requirements with the LEAST amount of changes to the application? (Select TWO.)

- A. Change the data model of the DynamoDB tables to ensure that all Scan and read operations meet DynamoDB best practices of uniform data access, reaching the full request throughput provisioned for the DynamoDB tables
- B. Enable DynamoDB auto scaling to manage the throughput capacity as table traffic increases Set the upper and lower limits to control costs and set a target utilization given the peak usage and how quickly the traffic changes.
- C. Provision Amazon ElastiCache for Redis with cluster mode enabled The cluster should be provisioned with enough shards to spread the application load and provision at least one read replica node for each shard
- D. Implement the DynamoDB Accelerator (DAX) client and provision a DAX cluster with the appropriate node types to sustain the application load
- E. Tune the item and query cache configuration for an optimal user experience
- F. Remove error retries and exponential backoffs in the application code to handle throttling errors

Answer: AE

NEW QUESTION 51

A company plans to move regulated and security-sensitive businesses to AWS. The Security team is developing a framework to validate the adoption of AWS best practice and industry-recognized compliance standards. The AWS Management Console is the preferred method for teams to provision resources. Which strategies should a Solutions Architect use to meet the business requirements and continuously assess, audit, and monitor the configurations of AWS resources? (Choose two.)

- A. Use AWS Config rules to periodically audit changes to AWS resources and monitor the compliance of the configuration
- B. Develop AWS Config custom rules using AWS Lambda to establish a test-driven development approach, and further automate the evaluation of configuration changes against the required controls.
- C. Use Amazon CloudWatch Logs agent to collect all the AWS SDK log
- D. Search the log data using a pre-defined set of filter patterns that machines mutating API call
- E. Send notifications using Amazon CloudWatch alarms when unintended changes are performed
- F. Archive log data by using a batch export to Amazon S3 and then Amazon Glacier for a long-term retention and auditability.
- G. Use AWS CloudTrail events to assess management activities of all AWS account
- H. Ensure that CloudTrail is enabled in all accounts and available AWS service
- I. Enable trails, encrypt CloudTrail event log files with an AWS KMS key, and monitor recorded activities with CloudWatch Logs.
- J. Use the Amazon CloudWatch Events near-real-time capabilities to monitor system events patterns, and trigger AWS Lambda functions to automatically revert non-authorized changes in AWS resource
- K. Also, target Amazon SNS topics to enable notifications and improve the response time of incident responses.
- L. Use CloudTrail integration with Amazon SNS to automatically notify unauthorized API activities. Ensure that CloudTrail is enabled in all accounts and available AWS service
- M. Evaluate the usage of Lambda functions to automatically revert non-authorized changes in AWS resources.

Answer: AC

Explanation:

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudwatch-alarms-for-cloudtrail.html>

https://docs.aws.amazon.com/en_pv/awscloudtrail/latest/userguide/best-practices-security.html

NEW QUESTION 52

The CISO of a large enterprise with multiple IT departments, each with its own AWS account, wants one central place where AWS permissions for users can be managed and users authentication credentials can be synchronized with the company's existing on-premises solution. Which solution will meet the CISO's requirements?

- A. Define AWS IAM roles based on the functional responsibilities of the users in a central account
- B. Create a SAML-based identity management provider
- C. Map users in the on-premises groups to IAM role
- D. Establish trust relationships between the other accounts and the central account.
- E. Deploy a common set of AWS IAM users, groups, roles, and policies in all of the AWS accounts using AWS Organization
- F. Implement federation between the on-premises identity provider and the AWS accounts.
- G. Use AWS Organizations in a centralized account to define service control policies (SCPs). Create a SAML-based identity management provider in each account and map users in the on-premises groups to AWS IAM roles.
- H. Perform a thorough analysis of the user base and create AWS IAM users accounts that have the necessary permission
- I. Set up a process to provision and de provision accounts based on data in the on-premises solution.

Answer: A

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 55

A Solution Architect is designing a deployment strategy for an application tier and has the following requirements.

- * The application code will need a 500 MB static dataset to be present before application startup.
- * The application tier be able to scale Up and down based on demand with as little startup time as possible.
- * The development team should be able to update the code multiple times each day.
- * Critical operating system (OS) patches must be installed within 48 hours of being released. Which deployment strategy meets these requirements?

- A. Use AWS Manager to create a new AMI with the updated OS patches . Update the Auto Scaling group to use the patches AMI and replace existing unpatched
- B. Use AWS CodeDeploy to push the application code to the instance
- C. Store the static data in Amazon EFS.
- D. Use AWS System Manager to create a new AMI with upload OS patches
- E. Update the Auto Scaling group to use the patches AMI and replace existing unpatches and the application code as a batch job every night
- F. Store the static data in Amazon EFS.
- G. Use an Amazon provided AMI for the OS Configure an Auto Scaling group set to a static instance count
- H. Configure an Amazon EC2 user data script to download the data from Amazon S3 install OS patches with AWS system Manager when they are released
- I. Use CodeDeploy to push the application code to the instances.
- J. Use an Amazon provided AMI for the OS Configure an Auto Scaling group Configure an Amazon EC2 user data script to download the data from Amazon S3. Replace existing instances after each Amazon-provided AMI releases

K. Use AWS CodeDeploy to push the application code to the instances.

Answer: C

NEW QUESTION 58

A Solutions Architect is working with a company that operates a standard three-tier web application in AWS. The web and application tiers run on Amazon EC2 and the database tier runs on Amazon RDS. The company is redesigning the web and application tiers to use Amazon API Gateway and AWS Lambda, and the company intends to deploy the new application within 6 months. The IT Manager has asked the Solutions Architect to reduce costs in the interim. Which solution will be MOST cost effective while maintaining reliability?

- A. Use Spot Instances for the web tier, On-Demand Instances for the application tier, and Reserved Instances for the database tier.
- B. Use On-Demand Instances for the web and application tiers, and Reserved Instances for the database tier.
- C. Use Spot Instances for the web and application tiers, and Reserved Instances for the database tier.
- D. Use Reserved Instances for the web, application, and database tiers.

Answer: B

NEW QUESTION 59

A financial services company logs personality identifiable information to its application logs stored in Amazon S3. Due to regulatory compliance requirements, the log files must be encrypted at rest. The Security team has mandated that the company's on-premises hardware security modules (HSMs) be used to generate the CMK material.

Which steps should the Solution Architect take to meet these requirements?

- A. Create an AWS CloudHSM cluster
- B. Create a new CMK in AWS KMS using AWS_CloudHSM as the source for the key material and an origin of AWS-CLOUDHSM
- C. Enable automatic key rotation on the CMK with a duration of 1 year
- D. Configure a bucket policy on the logging bucket to disallow uploads of unencrypted data and requires that the encryption source be AWS KMS.
- E. Provision an AWS Direct Connect connection, ensuring there is no overlap of the RFC 1918 address space between on-premises hardware and the VPC
- F. Configure an AWS bucket policy on the logging bucket requires all objects to be key material, and create a unique CMK for each logging event.
- G. Create a CMK in AWS KMS with no key material and an origin of EXTERNAL
- H. Import the key material generated from the on-premises HSMs into the CMK using the public key and import token provided by AWS
- I. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.
- J. Create a new CMK in AWS KMS with AWS-provided key material and an origin of AWS-KM
- K. Disable this CMK, and overwrite the key material with the material from the on-premises HSM using the public key and import token provided by AWS Re-enable the CMK
- L. Enable automatic key rotation on the CMK with a duration of 1 year
- M. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.

Answer: A

NEW QUESTION 62

A company is running a high-user-volume media-sharing application on premises. It currently hosts about 400 TB of data with millions of video files. The company is migrating this application to AWS to improve reliability and reduce costs.

The Solutions Architecture team plans to store the videos in an Amazon S3 bucket and use Amazon

CloudFront to distribute videos to users. The company needs to migrate this application to AWS within 10 days with the least amount of downtime possible. The company currently has 1 Gbps connectivity to the internet with 30 percent free capacity.

Which of the following solutions would enable the company to migrate the workload to AWS and meet all of the requirements?

- A. Use a multipart upload in Amazon S3 client to parallel-upload the data to the Amazon S3 bucket over the internet. Use the throttling feature to ensure that the Amazon S3 client does not use more than 30 percent of available internet capacity.
- B. Request an AWS Snowmobile with 1 PB capacity to be delivered to the data center. Load the data into Snowmobile and send it back to have AWS download that data to the Amazon S3 bucket. Sync the new data that was generated while migration was in flight.
- C. Use an Amazon S3 client to transfer data from the data center to the Amazon S3 bucket over the internet. Use the throttling feature to ensure the Amazon S3 client does not use more than 30 percent of available internet capacity.
- D. Request multiple AWS Snowball devices to be delivered to the data center. Load the data concurrently into these devices and send it back. Have AWS download that data to the Amazon S3 bucket. Sync the new data that was generated while migration was in flight.

Answer: D

Explanation:

<https://www.edureka.co/blog/aws-snowball-and-snowmobile-tutorial/>

NEW QUESTION 63

A company is using an Amazon CloudFront distribution to distribute both static and dynamic content from a web application running behind an Application Load Balancer. The web application requires user authorization and session tracking for dynamic content. The CloudFront distribution has a single cache behavior configured to forward the Authorization, Host, and User-Agent HTTP whitelist headers and a session cookie to the origin. All other cache behavior settings are set to their default value.

A valid ACM certificate is applied to the CloudFront distribution with a matching CNAME in the distribution settings. The ACM certificate is also applied to the HTTPS listener for the Application Load Balancer. The CloudFront origin protocol policy is set to HTTPS only. Analysis of the cache statistics report shows that the miss rate for this distribution is very high.

What can the Solutions Architect do to improve the cache hit rate for this distribution without causing the SSL/TLS handshake between CloudFront and the Application Load Balancer to fail?

- A. Create two cache behaviors for static and dynamic content
- B. Remove the User-Agent and Host HTTP headers from the whitelist headers section on both of the cache behaviors
- C. Remove the session cookie from the whitelist cookies section and the Authorization HTTP header from the whitelist headers section for cache behavior configured for static content.
- D. Remove the User-Agent and Authorization HTTP headers from the whitelist headers section of the cache behavior
- E. Then update the cache behavior to use presigned cookies for authorization.

- F. Remove the Host HTTP header from the whitelist headers section and remove the session cookie from the whitelist cookies section for the default cache behavior
- G. Enable automatic object compression and use Lambda@Edge viewer request events for user authorization.
- H. Create two cache behaviors for static and dynamic content
- I. Remove the User-Agent HTTP header from the whitelist headers section on both of the cache behavior
- J. Remove the session cookie from the whitelist cookies section and the Authorization HTTP header from the whitelist headers section for cache behavior configured for static content.

Answer: D

NEW QUESTION 68

A company wants to migrate its website from an on-premises data center onto AWS. At the same time, it wants to migrate the website to a containerized microservice-based architecture to improve the availability and cost efficiency. The company's security policy states that privileges and network permissions must be configured according to best practice, using least privilege.

A Solutions Architect must create a containerized architecture that meets the security requirements and has deployed the application to an Amazon ECS cluster. What steps are required after the deployment to meet the requirements? (Choose two.)

- A. Create tasks using the bridge network mode.
- B. Create tasks using the awsvpc network mode.
- C. Apply security groups to Amazon EC2 instances, and use IAM roles for EC2 instances to access other resources.
- D. Apply security groups to the tasks, and pass IAM credentials into the container at launch time to access other resources.
- E. Apply security groups to the tasks, and use IAM roles for tasks to access other resources.

Answer: BE

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2017/11/amazon-ecs-introduces-awsvpc-networking-mode-for-c>

<https://amazonaws-china.com/blogs/compute/introducing-cloud-native-networking-for-ecs-containers/>

<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-iam-roles.html>

NEW QUESTION 71

The Security team needs to provide a team of interns with an AWS environment so they can build the serverless video transcoding application. The project will use Amazon S3, AWS Lambda, Amazon API Gateway, Amazon Cognito, Amazon DynamoDB, and Amazon Elastic Transcoder.

The interns should be able to create and configure the necessary resources, but they may not have access to create or modify AWS IAM roles. The Solutions Architect creates a policy and attaches it to the interns' group.

How should the Security team configure the environment to ensure that the interns are self-sufficient?

- A. Create a policy that allows creation of project-related resources only
- B. Create roles with required service permissions, which are assumable by the services.
- C. Create a policy that allows creation of all project-related resources, including roles that allow access only to specified resources.
- D. Create roles with the required service permissions, which are assumable by the service
- E. Have the interns create and use a bastion host to create the project resources in the project subnet only.
- F. Create a policy that allows creation of project-related resources only
- G. Require the interns to raise a request for roles to be created with the Security team
- H. The interns will provide the requirements for the permissions to be set in the role.

Answer: A

NEW QUESTION 74

A company has a single AWS master billing account, which is the root of the AWS Organizations hierarchy. The company has multiple AWS accounts within this hierarchy, all organized into organization units (OUs). More OUs and AWS accounts will continue to be created as other parts of the business migrate applications to AWS. These business units may need to use different AWS services. The Security team is implementing the following requirements for all current and future AWS accounts.

* Control policies must be applied across all accounts to prohibit AWS servers.

* Exceptions to the control policies are allowed based on valid use cases. Which solution will meet these requirements with minimal optional overhead?

- A. Use an SCP in Organizations to implement a deny list of AWS server
- B. Apply this SCP at the level
- C. For any specific exceptions for an OU, create a new SCP for that OU and add the required AWS services to the allow list.
- D. Use an SCP in organizations to implement a deny list of AWS service
- E. Apply this SCP at the root level and each OU
- F. Remove the default AWS managed SCP from the root level and all OU level
- G. For any specific exceptions, modify the SCP attached to that OU, and add the required AWS required services to the allow list.
- H. Use an SCP in Organization to implement a deny list of AWS service
- I. Apply this SCP at each OU level
- J. Leave the default AWS managed SCP at the root level. For any specific exceptions for an OU, create a new SCP for that OU.
- K. Use an SCP in Organizations to implement an allow list of AWS service
- L. Apply this SCP at the root level
- M. Remove the default AWS managed SCP from the root level and all OU level
- N. For any specific exceptions for an OU, modify the SCP attached to that OU, and add the required AWS services to the allow list.

Answer: B

NEW QUESTION 79

A company has an application that generates a weather forecast that is updated every 15 minutes with an output resolution of 1 billion unique positions, each approximately 20 bytes in size (20 Gigabytes per forecast). Every hour, the forecast data is globally accessed approximately 5 million times (1,400 requests per second), and up to 10 times more during weather events. The forecast data is overwritten every update. Users of the current weather forecast application expect responses to queries to be returned in less than two seconds for each request.

Which design meets the required request rate and response time?

- A. Store forecast locations in an Amazon ES cluster
- B. Use an Amazon CloudFront distribution targeting an Amazon API Gateway endpoint with AWS Lambda functions responding to queries as the origin
- C. Enable API caching on the API Gateway stage with a cache-control timeout set for 15 minutes.
- D. Store forecast locations in an Amazon EFS volume
- E. Create an Amazon CloudFront distribution that targets an Elastic Load Balancing group of an Auto Scaling fleet of Amazon EC2 instances that have mounted the Amazon EFS volume
- F. Set the cache-control timeout for 15 minutes in the CloudFront distribution.
- G. Store forecast locations in an Amazon ES cluster
- H. Use an Amazon CloudFront distribution targeting an API Gateway endpoint with AWS Lambda functions responding to queries as the origin
- I. Create an Amazon Lambda@Edge function that caches the data locally at edge locations for 15 minutes.
- J. Store forecast locations in an Amazon S3 as individual object
- K. Create an Amazon CloudFront distribution targeting an Elastic Load Balancing group of an Auto Scaling fleet of EC2 instances, querying the origin of the S3 object
- L. Set the cache-control timeout for 15 minutes in the CloudFront distribution.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/lambdaedge-design-best-practices/>

NEW QUESTION 82

A company has an application written using an in-house software framework. The framework installation takes 30 minutes and is performed with a user data script. Company Developers deploy changes to the application frequently. The framework installation is becoming a bottleneck in this process. Which of the following would speed up this process?

- A. Create a pipeline to build a custom AMI with the framework installed and use this AMI as a baseline for application deployments.
- B. Employ a user data script to install the framework but compress the installation files to make them smaller.
- C. Create a pipeline to parallelize the installation tasks and call this pipeline from a user data script.
- D. Configure an AWS OpsWorks cookbook that installs the framework instead of employing user data
- E. Use this cookbook as a base for all deployments.

Answer: A

Explanation:

<https://aws.amazon.com/codepipeline/features/?nc=sn&loc=2>

NEW QUESTION 86

A company must deploy multiple independent instances of an application. The front-end application is internet accessible. However, corporate policy stipulates that the backends are to be isolated from each other and the internet, yet accessible from a centralized administration server. The application setup should be automated to minimize the opportunity for mistakes as new instances are deployed. Which option meets the requirements and MINIMIZES costs?

- A. Use an AWS CloudFormation template to create identical IAM roles for each region
- B. Use AWS CloudFormation StackSets to deploy each application instance by using parameters to customize for each instance, and use security groups to isolate each instance while permitting access to the central server.
- C. Create each instance of the application IAM roles and resources in separate accounts by using AWS CloudFormation StackSet
- D. Include a VPN connection to the VPN gateway of the central administration server.
- E. Duplicate the application IAM roles and resources in separate accounts by using a single CloudFormation template
- F. Include VPC peering to connect the VPC of each application instance to a central VPC.
- G. Use the parameters of the AWS CloudFormation template to customize the deployment into separate account
- H. Include a NAT gateway to allow communication back to the central administration server.

Answer: A

NEW QUESTION 91

A Solutions Architect is working with a company that is extremely sensitive to its IT costs and wishes to implement controls that will result in a predictable AWS spend each month.

Which combination of steps can help the company control and monitor its monthly AWS usage to achieve a cost that is as close as possible to the target amount? (Choose three.)

- A. Implement an IAM policy that requires users to specify a 'workload' tag for cost allocation when launching Amazon EC2 instances.
- B. Contact AWS Support and ask that they apply limits to the account so that users are not able to launch more than a certain number of instance types.
- C. Purchase all upfront Reserved Instances that cover 100% of the account's expected Amazon EC2 usage.
- D. Place conditions in the users' IAM policies that limit the number of instances they are able to launch.
- E. Define 'workload' as a cost allocation tag in the AWS Billing and Cost Management console.
- F. Set up AWS Budgets to alert and notify when a given workload is expected to exceed a defined cost.

Answer: AEF

NEW QUESTION 93

A company collects a steady stream of 10 million data records from 100,000 sources each day. These records are written to an Amazon RDS MySQL DB. A query must produce the daily average of a data source over the past 30 days. There are twice as many reads as writes. Queries to the collected data are for one source ID at a time.

How can the Solutions Architect improve the reliability and cost effectiveness of this solution?

- A. Use Amazon Aurora with MySQL in a Multi-AZ mode
- B. Use four additional read replicas.
- C. Use Amazon DynamoDB with the source ID as the partition key and the timestamp as the sort key
- D. Use a Time to Live (TTL) to delete data after 30 days.

- E. Use Amazon DynamoDB with the source ID as the partition key
- F. Use a different table each day.
- G. Ingest data into Amazon Kinesis using a retention period of 30 days
- H. Use AWS Lambda to write data records to Amazon ElastiCache for read access.

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html>

NEW QUESTION 96

A company has developed a web application that runs on Amazon EC2 instances in one AWS Region. The company has taken on new business in other countries and must deploy its application into other to meet low-latency requirements for its users. The regions can be segregated, and an application running in one region does not need to communicate with instances in other regions.

How should the company's Solutions Architect automate the deployment of the application so that it can be MOST efficiently deployed into multiple regions?

- A. Write a bash script that uses the AWS CLI to query the current state in one region and output a JSON representation
- B. Pass the JSON representation to the AWS CLI, specifying the --region parameter to deploy the application to other regions.
- C. Write a bash script that uses the AWS CLI to query the current state in one region and output an AWS CloudFormation template
- D. Create a CloudFormation stack from the template by using the AWS CLI, specifying the --region parameter to deploy the application to other regions.
- E. Write a CloudFormation template describing the application's infrastructure in the resources section. Create a CloudFormation stack from the template by using the AWS CLI, specify multiple regions using the --regions parameter to deploy the application.
- F. Write a CloudFormation template describing the application's infrastructure in the Resources section. Use a CloudFormation stack set from an administrator account to launch stack instances that deploy the application to other regions.

Answer: D

Explanation:

A stack set lets you create stacks in AWS accounts across regions by using a single AWS CloudFormation template. All the resources included in each stack are defined by the stack set's AWS CloudFormation template. As you create the stack set, you specify the template to use, as well as any parameters and capabilities that template requires. <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-concepts.html>

<https://sanderknape.com/2017/07/cloudformation-stacksets-automated-cross-account-region-deployments/>

NEW QUESTION 101

A financial company is using a high-performance compute cluster running on Amazon EC2 instances to perform market simulations. A DNS record must be created in an Amazon Route 53 private hosted zone when instances start. The DNS record must be removed after instances are terminated.

Currently the company uses a combination of Amazon CloudWatch Events and AWS Lambda to create the

DNS record. The solution worked well in testing with small clusters, but in production with clusters containing thousands of instances the company sees the following error in the Lambda logs:

HTTP 400 error (Bad request).

The response header also includes a status code element with a value of "Throttling" and a status message element with a value of "Rate exceeded "

Which combination of steps should the Solutions Architect take to resolve these issues? (Select THREE)

- A. Configure an Amazon SQS FIFO queue and configure a CloudWatch Events rule to use this queue as a target
- B. Remove the Lambda target from the CloudWatch Events rule
- C. Configure an Amazon Kinesis data stream and configure a CloudWatch Events rule to use this queue as a target. Remove the Lambda target from the CloudWatch Events rule
- D. Update the CloudWatch Events rule to trigger on Amazon EC2 "Instance Launch Successful" and "Instance Terminate Successful" events for the Auto Scaling group used by the cluster
- E. Configure a Lambda function to retrieve messages from an Amazon SQS queue. Modify the Lambda function to retrieve a maximum of 10 messages then batch the messages by Amazon Route 53 API call type and submit. Delete the messages from the SQS queue after successful API calls.
- F. Configure an Amazon SQS standard queue and configure the existing CloudWatch Events rule to use this queue as a target. Remove the Lambda target from the CloudWatch Events rule.
- G. Configure a Lambda function to read data from the Amazon Kinesis data stream and configure the batch window to 5 minutes. Modify the function to make a single API call to Amazon Route 53 with all records read from the Kinesis data stream

Answer: BEF

NEW QUESTION 104

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual SAP-C01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the SAP-C01 Product From:

<https://www.2passeasy.com/dumps/SAP-C01/>

Money Back Guarantee

SAP-C01 Practice Exam Features:

- * SAP-C01 Questions and Answers Updated Frequently
- * SAP-C01 Practice Questions Verified by Expert Senior Certified Staff
- * SAP-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * SAP-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year