

Amazon-Web-Services

Exam Questions SAP-C01

AWS Certified Solutions Architect- Professional



NEW QUESTION 1

A group of Amazon EC2 instances have been configured as high performance computing (HPC) cluster. The instances are running in a placement group, and are able to communicate with each other at network of up to 20 Gbps.

The cluster needs to communicate with a control EC2 instance outside of the placement group. The control instance has the same instance type and AMI as the other instances, and is configured with a public IP address.

How can the Solutions Architect improve the network speeds between the control instance and the instances in the placement group?

- A. Terminate the control instance and relaunch in the placement group.
- B. Ensure that the instances are communicating using the private IP addresses.
- C. Ensure that the control instance is using an Elastic Network Adapter.
- D. Move the control instance inside the placement group.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

NEW QUESTION 2

A company has an Amazon EC2 deployment that has the following architecture:



An application tier that contains 8 m4.xlarge instances

A Classic Load Balancer

Amazon S3 as a persistent data store

After one of the EC2 instances fails, users report very slow processing of their requests. A Solutions Architect must recommend design changes to maximize system reliability. The solution must minimize costs.

What should the Solution Architect recommend?

- A. Migrate the existing EC2 instances to a serverless deployment using AWS Lambda functions
- B. Change the Classic Load Balancer to an Application Load Balancer
- C. Replace the application tier with m4.large instances in an Auto Scaling group
- D. Replace the application tier with 4 m4.2xlarge instances

Answer: B

Explanation:

By default, connection draining is enabled for Application Load Balancers but must be enabled for Classic Load Balancers. When Connection Draining is enabled and configured, the process of deregistering an instance from an Elastic Load Balancer gains an additional step. For the duration of the configured timeout, the load balancer will allow existing, in-flight requests made to an instance to complete, but it will not send any new requests to the instance. During this time, the API will report the status of the instance as InService, along with a message stating that "Instance deregistration currently in progress." Once the timeout is reached, any remaining connections will be forcibly closed. <https://docs.aws.amazon.com/autoscaling/ec2/userguide/attach-load-balancer-asg.html>

<https://aws.amazon.com/blogs/aws/elb-connection-draining-remove-instances-from-service-with-care/>

NEW QUESTION 3

A company has a standard three-tier architecture using two Availability Zones. During the company's off season, users report that the website is not working. The Solutions Architect finds that no changes have been made to the environment recently, the website is reachable, and it is possible to log in. However, when the Solutions Architect selects the "find a store near you" function, the maps provided on the site by a third-party RESTful API call do not work about 50% of the time after refreshing the page. The outbound API calls are made through Amazon EC2 NAT instances.

What is the MOST likely reason for this failure and how can it be mitigated in the future?

- A. The network ACL for one subnet is blocking outbound web traffi
- B. Open the network ACL and prevent administration from making future changes through IAM.
- C. The fault is in the third-party environmen
- D. Contact the third party that provides the maps and request a fix that will provide better uptime.
- E. One NAT instance has become overloade
- F. Replace both EC2 NAT instances with a larger-sized instance and make sure to account for growth when making the new instance size.
- G. One of the NAT instances faile
- H. Recommend replacing the EC2 NAT instances with a NAT gateway.

Answer: D

Explanation:

The issue is 50% failure, means the balancing over 2 AZs is failing on one NAT instance in one AZ. The solution is to replace the NAT instance with fully managed and high available NAT gateway.

NEW QUESTION 4

A company is designing a new highly available web application on AWS. The application requires consistent and reliable connectivity from the application servers in AWS to a backend REST API hosted in the company's on-premises environment. The backend connection between AWS and on-premises will be routed over

an AWS Direct Connect connection through a private virtual interface. Amazon Route 53 will be used to manage private DNS records for the application to resolve the IP address on the backend REST API.

Which design would provide a reliable connection to the backend API?

- A. Implement at least two backend endpoints for the backend REST API, and use Route 53 health checks to monitor the availability of each backend endpoint and perform DNS-level failover.
- B. Install a second Direct Connect connection from a different network carrier and attach it to the same virtual private gateway as the first Direct Connect connection.
- C. Install a second cross connect for the same Direct Connect connection from the same network carrier, and join both connections to the same link aggregation group (LAG) on the same private virtual interface.
- D. Create an IPSec VPN connection routed over the public internet from the on-premises data center to AWS and attach it to the same virtual private gateway as the Direct Connect connection.

Answer: A

NEW QUESTION 5

A company uses an Amazon EMR cluster to process data once a day. The raw data comes from Amazon S3, and the resulting processed data is also stored in Amazon S3. The processing must complete within 4 hours; currently, it only takes 3 hours. However, the processing time is taking 5 to 10 minutes. longer each week due to an increasing volume of raw data.

The team is also concerned about rising costs as the compute capacity increases. The EMR cluster is currently running on three m3.xlarge instances (one master and two core nodes).

Which of the following solutions will reduce costs related to the increasing compute needs?

- A. Add additional task nodes, but have the team purchase an all-upfront convertible Reserved Instance for each additional node to offset the costs.
- B. Add additional task nodes, but use instance fleets with the master node in on-Demand mode and a mix of On-Demand and Spot Instances for the core and task node
- C. Purchase a scheduled Reserved Instances for the master node.
- D. Add additional task nodes, but use instance fleets with the master node in Spot mode and a mix of On-Demand and Spot Instances for the core and task node
- E. Purchase enough scheduled Reserved Instances to offset the cost of running any On-Demand instances.
- F. Add additional task nodes, but use instance fleets with the master node in On-Demand mode and a mix of On-Demand and Spot Instances for the core and task node
- G. Purchase a standard all-upfront Reserved Instance for the master node.

Answer: B

NEW QUESTION 6

A photo-sharing and publishing company receives 10,000 to 150,000 images daily. The company receives the images from multiple suppliers and users registered with the service. The company is moving to AWS and wants to enrich the existing metadata by adding data using Amazon Rekognition.

The following is an example of the additional data:

As part of the cloud migration program, the company uploaded existing image data to Amazon S3 and told users to upload images directly to Amazon S3. What should the Solutions Architect do to support these requirements?

- A. Trigger AWS Lambda based on an S3 event notification to create additional metadata using Amazon Rekognition
- B. Use Amazon DynamoDB to store the metadata and Amazon ES to create an index
- C. Use a web front-end to provide search capabilities backed by Amazon ES.
- D. Use Amazon Kinesis to stream data based on an S3 event
- E. Use an application running in Amazon EC2 to extract metadata from the image
- F. Then store the data on Amazon DynamoDB and Amazon CloudSearch and create an index
- G. Use a web front-end with search capabilities backed by CloudSearch.
- H. Start an Amazon SQS queue based on S3 event notification
- I. Then have Amazon SQS send the metadata information to Amazon DynamoDB
- J. An application running on Amazon EC2 extracts data from Amazon Rekognition using the API and adds data to DynamoDB and Amazon ES.
- K. Use a web front-end to provide search capabilities backed by Amazon ES.
- L. Trigger AWS Lambda based on an S3 event notification to create additional metadata using Amazon Rekognition
- M. Use Amazon RDS MySQL Multi-AZ to store the metadata information and use Lambda to create an index
- N. Use a web front-end with search capabilities backed by Lambda.

Answer: A

Explanation:

<https://github.com/aws-samples/lambda-refarch-imagerecognition>

NEW QUESTION 7

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 8

A company runs an application on a fleet of Amazon EC2 instances. The application requires low latency and random access to 100 GB of data. The application must be able to access the data at up to 3,000 IOPS. A Development team has configured the EC2 launch template to provision a 100-GB Provisioned IOPS (PIOPS) Amazon EBS volume with 3,000 IOPS provisioned. A Solutions Architect is tasked with lowering costs without impacting performance and durability. Which action should be taken?

- A. Create an Amazon EFS file system with the performance mode set to Max I/O. Configure the EC2 operating system to mount the EFS file system.
- B. Create an Amazon EFS file system with the throughput mode set to Provisioned. Configure the EC2 operating system to mount the EFS file system.
- C. Update the EC2 launch template to allocate a new 1-TB EBS General Purpose SSD (gp2) volume.
- D. Update the EC2 launch template to exclude the PIOPS volume. Configure the application to use local instance storage.

Answer: A

NEW QUESTION 9

A company would like to implement a serverless application by using Amazon API Gateway, AWS Lambda, and Amazon DynamoDB. They deployed a proof of concept and stated that the average response time is greater than what their upstream services can accept. Amazon CloudWatch metrics did not indicate any issues with DynamoDB but showed that some Lambda functions were hitting their timeout.

Which of the following actions should the Solutions Architect consider to improve performance? (Choose two.)

- A. Configure the AWS Lambda function to reuse containers to avoid unnecessary startup time.
- B. Increase the amount of memory and adjust the timeout on the Lambda function.
- C. Complete performance testing to identify the ideal memory and timeout configuration for the Lambda function.
- D. Create an Amazon ElastiCache cluster running Memcached, and configure the Lambda function for VPC integration with access to the Amazon ElastiCache cluster.
- E. Enable API cache on the appropriate stage in Amazon API Gateway, and override the TTL for individual methods that require a lower TTL than the entire stage.
- F. Increase the amount of CPU, and adjust the timeout on the Lambda function.
- G. Complete performance testing to identify the ideal CPU and timeout configuration for the Lambda function.

Answer: BD

Explanation:

<https://lumigo.io/blog/aws-lambda-timeout-best-practices/>

NEW QUESTION 10

A company has a 24 TB MySQL database in its on-premises data center that grows at the rate of 10 GB per day. The data center is connected to the company's AWS infrastructure with a 50 Mbps VPN connection.

The company is migrating the application and workload to AWS. The application code is already installed and tested on Amazon EC2. The company now needs to migrate the database and wants to go live on AWS within 3 weeks.

Which of the following approaches meets the schedule with LEAST downtime?

- A. 1. Use the VM Import/Export service to import a snapshot of the on-premises database into AWS. 2. Launch a new EC2 instance from the snapshot. 3. Set up ongoing database replication from on-premises to the EC2 database over the VPN. 4. Change the DNS entry to point to the EC2 database. 5. Stop the replication.
- B. 1. Launch an AWS DMS instance. 2. Launch an Amazon RDS Aurora MySQL DB instance. 3. Configure the AWS DMS instance with on-premises and Amazon RDS database information. 4. Start the replication task within AWS DMS over the VPN. 5. Change the DNS entry to point to the Amazon RDS MySQL database. 6. Stop the replication.
- C. 1. Create a database export locally using database-native tools. 2. Import that into AWS using AWS Snowball. 3. Launch an Amazon RDS Aurora DB instance. 4. Load the data in the RDS Aurora DB instance from the export. 5. Set up database replication from the on-premises database to the RDS Aurora DB instance over the VPN. 6. Change the DNS entry to point to the RDS Aurora DB instance. 7. Stop the replication.
- D. 1. Take the on-premises application offline. 2. Create a database export locally using database-native tools. 3. Import that into AWS using AWS Snowball. 4. Launch an Amazon RDS Aurora DB instance. 5. Load the data in the RDS Aurora DB instance from the export. 6. Change the DNS entry to point to the Amazon RDS Aurora DB instance. 7. Put the Amazon EC2 hosted application online.

Answer: C

NEW QUESTION 10

A large company has increased its utilization of AWS over time in an unmanaged way. As such, they have a large number of independent AWS accounts across different business units, projects, and environments. The company has created a Cloud Center of Excellence team, which is responsible for managing all aspects of the AWS Cloud, including their AWS accounts.

Which of the following should the Cloud Center of Excellence team do to BEST address their requirements in a centralized way? (Select two.)

- A. Control all AWS account root user credentials.
- B. Assign AWS IAM users in the account of each user who needs to access AWS resources.
- C. Follow the policy of least privilege in assigning permissions to each user.
- D. Tag all AWS resources with details about the business unit, project, and environment.
- E. Send all AWS Cost and Usage reports to a central Amazon S3 bucket, and use tools such as Amazon Athena and Amazon QuickSight to collect billing details by business unit.
- F. Use the AWS Marketplace to choose and deploy a Cost Management tool.
- G. Tag all AWS resources with details about the business unit, project, and environment.
- H. Send all AWS Cost and Usage reports for the AWS accounts to this tool for analysis.
- I. Set up AWS Organization.
- J. Enable consolidated billing, and link all existing AWS accounts to a master billing account.
- K. Tag all AWS resources with details about the business unit, project, and environment.
- L. Analyze Cost and Usage reports using tools such as Amazon Athena and Amazon QuickSight to collect billing details by business unit.
- M. Using a master AWS account, create IAM users within the master account.
- N. Define IAM roles in the other AWS accounts, which cover each of the required functions in the account.
- O. Follow the policy of least privilege in assigning permissions to each role, then enable the IAM users to assume the roles that they need to use.

Answer: DE

NEW QUESTION 11

A Solutions Architect is responsible for redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application runs on a single high-memory Amazon EC2 instance. It accepts HTTP requests from upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average., most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel.

Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backing. In addition, the current system has issues with availability and data if the single application node fails.

Clients that access this service cannot be modified. They expect to receive a response to each HTTP request they send within 10 seconds before they will time out and retry the request.

Which approach would improve the availability and durability of the system while decreasing the processing latency and minimizing costs?

- A. Create an Amazon API Gateway REST API that uses Lambda proxy integration to pass requests to an AWS Lambda function
- B. Migrate the core processing code to a Lambda function and write a wrapper class that provides a handler method that converts the proxy events to the internal application data model and invokes the processing module.
- C. Create an Amazon API Gateway REST API that uses a service proxy to put items in an Amazon SQS queue
- D. Extract the core processing code from the existing application and update it to pull items from Amazon SQS queue
- E. Extract the core processing code from the existing application and update it to pull items from Amazon SQS instead of an in-memory queue
- F. Deploy the new processing application to smaller EC2 instances within an Auto Scaling group that scales dynamically based on the approximate number of messages in the Amazon SQS queue.
- G. Modify the application to use Amazon DynamoDB instead of Amazon RD
- H. Configure Auto Scaling for the DynamoDB tabl
- I. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilizatio
- J. Back the in-memory queue with a memory-mapped file to an instance store volume and periodically write that file to Amazon S3.
- K. Update the application to use a Redis task queue instead of the in-memory queue
- L. Build a Docker container image for the applicatio
- M. Create an Amazon ECS task definition that includes the application container and a separate container to host Redi
- N. Deploy the new task definition as an ECS service using AWS Fargate and enable Auto Scaling.

Answer: B

NEW QUESTION 16

During an audit a Security team discovered that a Development team was putting IAM user secret access keys in their code and then committing it to an AWS CodeCommit repository The Security team wants to automatically find and remediate instances of this security vulnerability

Which solution will ensure that the credentials are appropriately secured automatically?

- A. Run a script nightly using AWS Systems Manager Run Command to search (or credentials on the development instances It found, use AWS Secrets Manager to rotate the credentials
- B. Use a scheduled AWS Lambda function to download and scan the application code from CodeCommit If credentials are found generate new credentials and store them in AWS KMS
- C. Configure Amazon Macie to scan for credentials in CodeCommit repositories If credentials are found, trigger an AWS Lambda function to disable the credentials and notify the user
- D. Configure a CodeCommit trigger to invoke an AWS Lambda function to scan new code submissions for credentials If credentials are found, disable them in AWS IAM and notify the user

Answer: C

NEW QUESTION 17

A company uses Amazon S3 to store documents that may only be accessible to an Amazon EC2 instance in a certain virtual private cloud (VPC). The company fears that a malicious insider with access to this instance could also set up an EC2 instance in another VPC to access these documents.

Which of the following solutions will provide the required protection?

- A. Use an S3 VPC endpoint and an S3 bucket policy to limit access to this VPC endpoint.
- B. Use EC2 instance profiles and an S3 bucket policy to limit access to the role attached to the instance profile.
- C. Use S3 client-side encryption and store the key in the instance metadata.
- D. Use S3 server-side encryption and protect the key with an encryption context.

Answer: A

Explanation:

<https://docs.aws.amazon.com/vpc/latest/userguide/vpce-gateway.html>

Endpoint connections cannot be extended out of a VPC. Resources on the other side of a VPN connection, VPC peering connection, AWS Direct Connect connection, or ClassicLink connection in your VPC cannot use the endpoint to communicate with resources in the endpoint service.

NEW QUESTION 22

A company is moving a business-critical, multi-tier application to AWS. The architecture consists of a desktop client application and server infrastructure. The server infrastructure resides in an on-premises data center that frequently fails to maintain the application uptime SLA of 99.95%. A Solutions Architect must re-architect the application to ensure that it can meet or exceed the SLA.

The application contains a PostgreSQL database running on a single virtual machine. The business logic and presentation layers are load balanced between multiple virtual machines. Remote users complain about slow load times while using this latency-sensitive application.

Which of the following will meet the availability requirements with little change to the application while improving user experience and minimizing costs?

- A. Migrate the database to a PostgreSQL database in Amazon EC2. Host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balance
- B. Allocate an Amazon WorkSpaces Workspace for each end user to improve the user experience.
- C. Migrate the database to an Amazon RDS Aurora PostgreSQL configuratio
- D. Host the application and presentation layers in an Auto Scaling configuration on Amazon EC2 instances behind an Application Load Balance

- E. Use Amazon AppStream 2.0 to improve the user experience.
- F. Migrate the database to an Amazon RDS PostgreSQL Multi-AZ configuratio
- G. Host the application and presentation layers in automatically scaled AWS Fargate containers behind a Network Load Balance
- H. Use Amazon ElastiCache to improve the user experience.
- I. Migrate the database to an Amazon Redshift cluster with at least two node
- J. Combine and host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balance
- K. Use Amazon CloudFront to improve the user experience.

Answer: B

NEW QUESTION 23

A company has developed a new billing application that will be released in two weeks. Developers are testing the application running on 10 EC2 instances managed by an Auto Scaling group in subnet 172.31.0.0/24 within VPC A with CIDR block 172.31.0.0/16. The Developers noticed connection timeout errors in the application logs while connecting to an Oracle database running on an Amazon EC2 instance in the same region within VPC B with CIDR block 172.50.0.0/16. The IP of the database instance is hard-coded in the application instances.

Which recommendations should a Solutions Architect present to the Developers to solve the problem in a secure way with minimal maintenance and overhead?

- A. Disable the SrcDestCheck attribute for all instances running the application and Oracle Database. Change the default route of VPC A to point ENI of the Oracle Database that has an IP address assigned within the range of 172.50.0.0/26
- B. Create and attach internet gateways for both VPC
- C. Configure default routes to the Internet gateways for both VPC
- D. Assign an Elastic IP for each Amazon EC2 instance in VPC A
- E. Create a VPC peering connection between the two VPCs and add a route to the routing table of VPC A that points to the IP address range of 172.50.0.0/16
- F. Create an additional Amazon EC2 instance for each VPC as a customer gateway; create one virtual private gateway (VGW) for each VPC, configure an end-to-end VPC, and advertise the routes for 172.50.0.0/16

Answer: C

NEW QUESTION 26

A bank is designing an online customer service portal where customers can chat with customer service agents. The portal is required to maintain a 15-minute RPO or RTO in case of a regional disaster. Banking regulations require that all customer service chat transcripts must be preserved on durable storage for at least 7 years, chat conversations must be encrypted in-flight, and transcripts must be encrypted at rest. The Data Lost Prevention team requires that data at rest must be encrypted using a key that the team controls, rotates, and revokes.

Which design meets these requirements?

- A. The chat application logs each chat message into Amazon CloudWatch Log
- B. A scheduled AWS Lambda function invokes a CloudWatch Log
- C. CreateExportTask every 5 minutes to export chat transcripts to Amazon S3. The S3 bucket is configured for cross-region replication to the backup regio
- D. Separate AWS KMS keys are specified for the CloudWatch Logs group and the S3 bucket.
- E. The chat application logs each chat message into two different Amazon CloudWatch Logs groups in two different regions, with the same AWS KMS key applie
- F. Both CloudWatch Logs groups are configured to export logs into an Amazon Glacier vault with a 7-year vault lock policy with a KMS key specified.
- G. The chat application logs each chat message into Amazon CloudWatch Log
- H. A subscription filter on the CloudWatch Logs group feeds into an Amazon Kinesis Data Firehose which streams the chat messages into an Amazon S3 bucket in the backup regio
- I. Separate AWS KMS keys are specified for the CloudWatch Logs group and the Kinesis Data Firehose.
- J. The chat application logs each chat message into Amazon CloudWatch Log
- K. The CloudWatch Logs group is configured to export logs into an Amazon Glacier vault with a 7-year vault lock polic
- L. Glacier cross-region replication mirrors chat archives to the backup regio
- M. Separate AWS KMS keys are specified for the CloudWatch Logs group and the Amazon Glacier vault.

Answer: B

NEW QUESTION 29

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 deman
- C. Save AWS Simple Monthly Calculator reports in Amazon S3 for trend analysi
- 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 AWSWell-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 budge
- L. Create Amazon EC2 Auto Scaling groups for applications that experience fluctuating deman
- 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 spendin
- O. Use Spot instances on nightly batch processing jobs.

Answer: B

Explanation:

Import/Export supports importing and exporting data into and out of Amazon S3 buckets. For significant data sets, AWS Import/Export is often faster than Internet

transfer and more cost effective than upgrading your connectivity.

NEW QUESTION 34

A company is running a .NET three-tier web application on AWS. The team currently uses XL storage optimized instances to store serve the website's image and video files on local instance storage. The company has encountered issues with data loss from replication and instance failures. The Solutions Architect has been asked to redesign this application to improve its reliability while keeping costs low.

Which solution will meet these requirements?

- A. Set up a new Amazon EFS share, move all image and video files to this share, and then attach this new drive as a mount point to all existing server
- B. Create an Elastic Load Balancer with Auto Scaling general purpose instance
- C. Enable Amazon CloudFront to the Elastic Load Balance
- D. Enable Cost Explorer and use AWS Trusted advisor checks to continue monitoring the environment for future savings.
- E. Implement Auto Scaling with general purpose instance types and an Elastic Load Balance
- F. Enable an Amazon CloudFront distribution to Amazon S3 and move images and video files to Amazon S3. Reserve general purpose instances to meet base performance requirement
- G. Use Cost Explorer and AWS Trusted Advisor checks to continue monitoring the environment for future savings.
- H. Move the entire website to Amazon S3 using the S3 website hosting feature
- I. Remove all the web servers and have Amazon S3 communicate directly with the application servers in Amazon VPC.
- J. Use AWS Elastic Beanstalk to deploy the .NET application
- K. Move all images and video files to Amazon EFS
- L. Create an Amazon CloudFront distribution that points to the EFS share
- M. Reserve the m4.xl instances needed to meet base performance requirements.

Answer: B

NEW QUESTION 35

A company has an Amazon VPC that is divided into a public subnet and a private subnet. A web application runs in Amazon VPC, and each subnet has its own NACL. The public subnet has a CIDR of 10.0.0.0/24. An Application Load Balancer is deployed to the public subnet. The private subnet has a CIDR of 10.0.1.0/24. Amazon EC2 instances that run a web server on port 80 are launched into the private subnet.

Only network traffic that is required for the Application Load Balancer to access the web application can be allowed to travel between the public and private subnets.

What collection of rules should be written to ensure that the private subnet's NACL meets the requirement? (Select TWO.)

- A. An inbound rule for port 80 from source 0.0.0.0/0
- B. An inbound rule for port 80 from source 10.0.0.0/24
- C. An outbound rule for port 80 to destination 0.0.0.0/0
- D. An outbound rule for port 80 to destination 10.0.0.0/24
- E. An outbound rule for ports 1024 through 65535 to destination 10.0.0.0/24

Answer: BE

NEW QUESTION 39

A Solutions Architect is designing a network solution for a company that has applications running in a data center in Northern Virginia. The applications in the company's data center require predictable performance to applications running in a virtual private cloud (VPC) located in us-east-1, and a secondary VPC in us-west-2 within the same account. The company data center is collocated in an AWS Direct Connect facility that serves the us-east-1 region. The company has already ordered an AWS Direct Connect connection and a cross-connect has been established.

Which solution will meet the requirements at the LOWEST cost?

- A. Provision a Direct Connect gateway and attach the virtual private (VGW) for the VPC in us-east-1 and the VGW for the VPC in us-west-2. Create a private VIF on the Direct Connect connection and associate it to the Direct Connect gateway.
- B. Create private VIFs on the Direct Connect connection for each of the company's VPCs in the us-east-1 and us-west-2 region.
- C. Configure the company's data center router to connect directly with the VPCs in those regions via the private VIFs.
- D. Deploy a transit VPC solution using Amazon EC2-based router instances in the us-east-1 region. Establish IPsec VPN tunnels between the transit routers and virtual private gateways (VGWs) located in the us-east-1 and us-west-2 regions, which are attached to the company's VPCs in those regions.
- E. Create a public VIF on the Direct Connect connection and establish IPsec VPN tunnels over the public VIF between the transit routers and the company's data center router.
- F. Order a second Direct Connect connection to a Direct Connect facility with connectivity to the us-west-2 region.
- G. Work with partner to establish a network extension link over dark fiber from the Direct Connect facility to the company's data center.
- H. Establish private VIFs on the Direct Connect connections for each of the company's VPCs in the respective region.
- I. Configure the company's data center router to connect directly with the VPCs in those regions via the private VIFs.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/aws/new-aws-direct-connect-gateway-inter-region-vpc-access/>

NEW QUESTION 40

A company stores sales transaction data in Amazon DynamoDB tables. To detect anomalous behaviors and respond quickly, all changes to the items stored in the DynamoDB tables must be logged within 30 minutes. Which solution meets the requirements?

- A. Copy the DynamoDB tables into Apache Hive tables on Amazon EMR every hour and analyze them for anomalous behavior.
- B. Send Amazon SNS notifications when anomalous behaviors are detected.
- C. Use AWS CloudTrail to capture all the APIs that change the DynamoDB table.
- D. Send SNS notifications when anomalous behaviors are detected using CloudTrail event filtering.
- E. Use Amazon DynamoDB Streams to capture and send updates to AWS Lambda.
- F. Create a Lambda function to output records to Amazon Kinesis Data Stream.
- G. Analyze any anomalies with Amazon Kinesis Data Analytics.
- H. Send SNS notifications when anomalous behaviors are detected.
- I. Use event patterns in Amazon CloudWatch Events to capture DynamoDB API call events with an AWS Lambda function as a target to analyze behavior.

J. Send SNS notifications when anomalous behaviors are detected.

Answer: C

Explanation:

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

NEW QUESTION 44

A company wants to ensure that the workloads for each of its business units have complete autonomy and a minimal blast radius in AWS. The Security team must be able to control access to the resources and services in the account to ensure that particular services are not used by the business units.

How can a Solutions Architect achieve the isolation requirements?

- A. Create individual accounts for each business unit and add the account to an OU in AWS Organizations. Modify the OU to ensure that the particular services are blocked.
- B. Federate each account with an IdP, and create separate roles for the business units and the Security team.
- C. Create individual accounts for each business unit.
- D. Federate each account with an IdP and create separate roles and policies for business units and the Security team.
- E. Create one shared account for the entire company.
- F. Create separate VPCs for each business unit.
- G. Create individual IAM policies and resource tags for each business unit.
- H. Federate each account with an IdP, and create separate roles for the business units and the Security team.
- I. Create one shared account for the entire company.
- J. Create individual IAM policies and resource tags for each business unit.
- K. Federate the account with an IdP, and create separate roles for the business units and the Security team.

Answer: A

NEW QUESTION 48

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 deployments.
- 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 stacks.
- 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 49

A company has an application that uses Amazon EC2 instances in an Auto Scaling group. The Quality Assurance (QA) department needs to launch a large number of short-lived environments to test the application. The application environments are currently launched by the Manager of the department using an AWS CloudFormation template. To launch the stack, the Manager uses a role with permission to use CloudFormation, EC2 and Auto Scaling APIs. The Manager wants to allow testers to launch their own environments, but does not want to grant broad permission to each user. Which set up would achieve these goals?

- A. Upload the AWS CloudFormation template to Amazon S3. Give users in the QA department permission to assume the Manager's role and add a policy that restricts the permissions to the template and the resources it creates.
- B. Train users to launch the template from the CloudFormation console.
- C. Create an AWS Service Catalog product from the environment template.
- D. Add a launch constraint to the product with the existing role.
- E. Give users in the QA department permission to use AWS Service Catalog APIs only.
- F. Train users to launch the templates from the AWS Service Catalog console.
- G. Upload the AWS CloudFormation template to Amazon S3. Give users in the QA department permission to use CloudFormation and S3 APIs, with conditions that restrict the permission to the template and the resources it creates.
- H. Train users to launch the template from the CloudFormation console.
- I. Create an AWS Elastic Beanstalk application from the environment template.
- J. Give users in the QA department permission to use Elastic Beanstalk permissions only.
- K. Train users to launch Elastic Beanstalk environments with the Elastic Beanstalk CLI, passing the existing role to the environment as a service role.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/mt/how-to-launch-secure-and-governed-aws-resources-with-aws-cloudformation->

NEW QUESTION 52

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 region
- 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 region
- 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 region
- J. Use Amazon Route 53 to automatically switch traffic from one region to another using health check routing policies.

Answer: C

NEW QUESTION 53

A company used Amazon EC2 instances to deploy a web fleet to host a blog site. The EC2 instances are behind an Application Load Balancer (ALB) and are configured in an Auto Scaling group. The web application stores all blog content on an Amazon EFS volume.

The company recently added a feature for bloggers to add video to their posts, attracting 10 times the previous user traffic. At peak times of day, users report buffering and timeout issues while attempting to reach the site or watch videos.

Which is the MOST cost-efficient and scalable deployment that will resolve the issues for users?

- A. Reconfigure Amazon EFS to enable maximum I/O.
- B. Update the blog site to use instance store volumes for storag
- C. Copy the site contents to the volumes at launch and to Amazon S3 at shutdown.
- D. Configure an Amazon CloudFront distributio
- E. Point the distribution to an S3 bucket, and migrate the videos from EFS to Amazon S3.
- F. Set up an Amazon CloudFront distribution for all suite contents, and point the distribution at the ALB.

Answer: C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-https-connection-fails/>

NEW QUESTION 56

A Solutions Architect is designing the storage layer for a recently purchased application. The application will be running on Amazon EC2 instances and has the following layers and requirements:

Data layer: A POSIX file system shared across many systems.

Service layer: Static file content that requires block storage with more than 100k IOPS. Which combination of AWS services will meet these needs? (Choose two.)

- A. Data layer – Amazon S3
- B. Data layer – Amazon EC2 Ephemeral Storage
- C. Data layer – Amazon EFS
- D. Service layer – Amazon EBS volumes with Provisioned IOPS
- E. Service layer – Amazon EC2 Ephemeral Storage

Answer: CE

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/storage-optimized-instances.html>

NEW QUESTION 57

A Solutions Architect must update an application environment within AWS Elastic Beanstalk using a blue/green deployment methodology. The Solutions Architect creates an environment that is identical to the existing application environment and deploys the application to the new environment.

What should be done next to complete the update?

- A. Redirect to the new environment using Amazon Route 53
- B. Select the Swap Environment URLs option
- C. Replace the Auto Scaling launch configuration
- D. Update the DNS records to point to the green environment

Answer: B

Explanation:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAMESwap.html>

NEW QUESTION 61

A company runs an IoT platform on AWS. IoT sensors in various locations send data to the company's Node.js API servers on Amazon EC2 instances running behind an Application Load Balancer. The data is stored in an Amazon RDS MySQL DB instance that uses a 4 TB General Purpose SSD volume.

The number of sensors the company has deployed in the field has increased over time, and is expected to grow significantly. The API servers are consistently

overloaded and RDS metrics show high write latency.

Which of the following steps together will resolve the issues permanently and enable growth as new sensors are provisioned, while keeping this platform cost-efficient? (Choose two.)

- A. Resize the MySQL General Purpose SSD storage to 6 TB to improve the volume's IOPS
- B. Re-architect the database tier to use Amazon Aurora instead of an RDS MySQL DB instance and add read replicas
- C. Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data
- D. Use AWS-X-Ray to analyze and debug application issues and add more API servers to match the load
- E. Re-architect the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance

Answer: CE

NEW QUESTION 65

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 instance
- B. Create a script that will check and repair the database upon reboot
- 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 two
- 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 a minimum instance count of one
- 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 load
- 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 69

What combination of steps could a Solutions Architect take to protect a web workload running on Amazon EC2 from DDoS and application layer attacks? (Select two.)

- A. Put the EC2 instances behind a Network Load Balancer and configure AWS WAF on it.
- B. Migrate the DNS to Amazon Route 53 and use AWS Shield
- C. Put the EC2 instances in an Auto Scaling group and configure AWS WAF on it.
- D. Create and use an Amazon CloudFront distribution and configure AWS WAF on it.
- E. Create and use an internet gateway in the VPC and use AWS Shield.

Answer: BD

Explanation:

References: <https://aws.amazon.com/answers/networking/aws-ddos-attack-mitigation/>

NEW QUESTION 74

A media company has a 30-TB repository of digital news videos. These videos are stored on tape in an on-premises tape library and referenced by a Media Asset Management (MAM) system. The company wants to enrich the metadata for these videos in an automated fashion and put them into a searchable catalog by using a MAM feature. The company must be able to search based on information in the video, such as objects, scenery items, or people's faces. A catalog is available that contains faces of people who have appeared in the videos that include an image of each person. The company would like to migrate these videos to AWS.

The company has a high-speed AWS Direct Connect connection with AWS and would like to move the MAM solution video content directly from its current file system.

How can these requirements be met by using the LEAST amount of ongoing management overhead and causing MINIMAL disruption to the existing system?

- A. Set up an AWS Storage Gateway, file gateway appliance on-premise
- B. Use the MAM solution to extract the videos from the current archive and push them into the file gateway
- C. Use the catalog of faces to build a collection in Amazon Rekognition
- D. Build an AWS Lambda function that invokes the Rekognition JavaScript SDK to have Rekognition pull the video from the Amazon S3 files backing the file gateway, retrieve the required metadata, and push the metadata into the MAM solution.
- E. Set up an AWS Storage Gateway, tape gateway appliance on-premise
- F. Use the MAM solution to extract the videos from the current archive and push them into the tape gateway
- G. Use the catalog of faces to build a collection in Amazon Rekognition
- H. Build an AWS Lambda function that invokes the Rekognition JavaScript SDK to have Amazon Rekognition process the video in the tape gateway, retrieve the required metadata, and push the metadata into the MAM solution.
- I. Configure a video ingestion stream by using Amazon Kinesis Video Stream
- J. Use the catalog of faces to build a collection in Amazon Rekognition
- K. Stream the videos from the MAM solution into Kinesis Video Stream
- L. Configure Amazon Rekognition to process the streamed video
- M. Then, use a stream consumer to retrieve the required metadata, and push the metadata into the MAM solution

- N. Configure the stream to store the videos in Amazon S3.
- O. Set up an Amazon EC2 instance that runs the OpenCV library
- P. Copy the videos, images, and face catalog from the on-premises library into an Amazon EBS volume mounted on this EC2 instance
- Q. Process the videos to retrieve the required metadata, and push the metadata into the MAM solution while also copying the video files to an Amazon S3 bucket.

Answer: C

Explanation:

<https://docs.aws.amazon.com/rekognition/latest/dg/streaming-video.html>

NEW QUESTION 76

A group of research institutions and hospitals are in a partnership to study 2 PBs of genomic data. The institute that owns the data stores it in an Amazon S3 bucket and updates it regularly. The institute would like to give all of the organizations in the partnership read access to the data. All members of the partnership are extremely cost-conscious, and the institute that owns the account with the S3 bucket is concerned about covering the costs for requests and data transfers from Amazon S3.

Which solution allows for secure datasharing without causing the institute that owns the bucket to assume all the costs for S3 requests and data transfers?

- A. Ensure that all organizations in the partnership have AWS account
- B. In the account with the S3 bucket, create a cross-account role for each account in the partnership that allows read access to the data
- C. Have the organizations assume and use that read role when accessing the data.
- D. Ensure that all organizations in the partnership have AWS account
- E. Create a bucket policy on the bucket that owns the data
- F. The policy should allow the accounts in the partnership read access to the bucket
- G. Enable Requester Pays on the bucket
- H. Have the organizations use their AWS credentials when accessing the data.
- I. Ensure that all organizations in the partnership have AWS account
- J. Configure buckets in each of the accounts with a bucket policy that allows the institute that owns the data the ability to write to the bucket
- K. Periodically sync the data from the institute's account to the other organization
- L. Have the organizations use their AWS credentials when accessing the data using their accounts.
- M. Ensure that all organizations in the partnership have AWS account
- N. In the account with the S3 bucket, create a cross-account role for each account in the partnership that allows read access to the data
- O. Enable Requester Pays on the bucket
- P. Have the organizations assume and use that read role when accessing the data.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/dev/RequesterPaysBuckets.html>

NEW QUESTION 77

A company wants to allow its Marketing team to perform SQL queries on customer records to identify market segments. The data is spread across hundreds of files. The records must be encrypted in transit and at rest. The Team Manager must have the ability to manage users and groups, but no team members should have access to services or resources not required for the SQL queries. Additionally, Administrators need to audit the queries made and receive notifications when a query violates rules defined by the Security team.

AWS Organizations has been used to create a new account and an AWS IAM user with administrator permissions for the Team Manager.

Which design meets these requirements?

- A. Apply a service control policy (SCP) that allows access to IAM, Amazon RDS, and AWS CloudTrail. Load customer records in Amazon RDS MySQL and train users to execute queries using the AWS CLI
- B. Stream the query logs to Amazon CloudWatch Logs from the RDS database instance
- C. Use a subscription filter with AWS Lambda functions to audit and alarm on queries against personal data.
- D. Apply a service control policy (SCP) that denies access to all services except IAM, Amazon Athena, Amazon S3, and AWS CloudTrail
- E. Store customer record files in Amazon S3 and train users to execute queries using the CLI via Athena
- F. Analyze CloudTrail events to audit and alarm on queries against personal data.
- G. Apply a service control policy (SCP) that denies access to all services except IAM, Amazon DynamoDB, and AWS CloudTrail
- H. Store customer records in DynamoDB and train users to execute queries using the AWS CLI
- I. Enable DynamoDB streams to track the queries that are issued and use an AWS Lambda function for real-time monitoring and alerting.
- J. Apply a service control policy (SCP) that allows access to IAM, Amazon Athena, Amazon S3, and AWS CloudTrail
- K. Store customer records as files in Amazon S3 and train users to leverage the Amazon S3 Select feature and execute queries using the AWS CLI
- L. Enable S3 object-level logging and analyze CloudTrail events to audit and alarm on queries against personal data.

Answer: D

NEW QUESTION 82

The Solutions Architect manages a serverless application that consists of multiple API gateways, AWS Lambda functions, Amazon S3 buckets, and Amazon DynamoDB tables. Customers say that a few application components slow while loading dynamic images, and some are timing out with the "504 Gateway Timeout" error. While troubleshooting the scenario, the Solutions Architect confirms that DynamoDB monitoring metrics are at acceptable levels.

Which of the following steps would be optimal for debugging these application issues? (Choose two.)

- A. Parse HTTP logs in Amazon API Gateway for HTTP errors to determine the root cause of the errors.
- B. Parse Amazon CloudWatch Logs to determine processing times for requested images at specified intervals.
- C. Parse VPC Flow Logs to determine if there is packet loss between the Lambda function and S3.
- D. Parse AWS X-Ray traces and analyze HTTP methods to determine the root cause of the HTTP errors.
- E. Parse S3 access logs to determine if objects being accessed are from specific IP addresses to narrow the scope to geographic latency issues.

Answer: BD

Explanation:

Firstly "A 504 Gateway Timeout Error means your web server didn't receive a timely response from another server upstream when it attempted to load one of your web pages. Put simply, your web servers aren't communicating with each other fast enough". This specific issue is addressed in the AWS article "Tracing, Logging

and Monitoring an API Gateway API". https://docs.amazonaws.cn/en_us/apigateway/latest/developerguide/monitoring_overview.html

NEW QUESTION 87

A company has a High Performance Computing (HPC) cluster in its on-premises data center which runs thousands of jobs in parallel for one week every month, processing petabytes of images. The images are stored on a network file server, which is replicated to a disaster recovery site. The on-premises data center has reached capacity and has started to spread the jobs out over the course of month in order to better utilize the cluster, causing a delay in the job completion. The company has asked its Solutions Architect to design a cost-effective solution on AWS to scale beyond the current capacity of 5,000 cores and 10 petabytes of data. The solution must require the least amount of management overhead and maintain the current level of durability. Which solution will meet the company's requirements?

- A. Create a container in the Amazon Elastic Container Registry with the executable file for the job
- B. Use Amazon ECS with Spot Fleet in Auto Scaling group
- C. Store the raw data in Amazon EBS SC1 volumes and write the output to Amazon S3.
- D. Create an Amazon EMR cluster with a combination of On Demand and Reserved Instance Task Nodes that will use Spark to pull data from Amazon S3. Use Amazon DynamoDB to maintain a list of jobs that need to be processed by the Amazon EMR cluster.
- E. Store the raw data in Amazon S3, and use AWS Batch with Managed Compute Environments to create Spot Fleet
- F. Submit jobs to AWS Batch Job Queues to pull down objects from Amazon S3 onto Amazon EBS volumes for temporary storage to be processed, and then write the results back to Amazon S3.
- G. Submit the list of jobs to be processed to an Amazon SQS to queue the jobs that need to be processed. Create a diversified cluster of Amazon EC2 worker instances using Spot Fleet that will automatically scale based on the queue depth
- H. Use Amazon EFS to store all the data sharing it across all instances in the cluster.

Answer: B

NEW QUESTION 91

A company has been using a third-party provider for its content delivery network and recently decided to switch to Amazon CloudFront. The Development team wants to maximize performance for the global user base. The company uses a content management system (CMS) that serves both static and dynamic content. The CMS is behind an Application Load Balancer (ALB) which is set as the default origin for the distribution. Static assets are served from an Amazon S3 bucket. The Origin Access Identity (OAI) was created properly and the S3 bucket policy has been updated to allow the GetObject action from the OAI, but static assets are receiving a 404 error. Which combination of steps should the Solutions Architect take to fix the error? (Select TWO.)

- A. Add another origin to the CloudFront distribution for the static assets
- B. Add a path-based rule to the ALB to forward requests for the static assets
- C. Add an RTMP distribution to allow caching of both static and dynamic content
- D. Add a behavior to the CloudFront distribution for the path pattern and the origin of the static assets
- E. Add a host header condition to the ALB listener and forward the header from CloudFront to add traffic to the allow list

Answer: AD

NEW QUESTION 92

A Solutions Architect is migrating a 10 TB PostgreSQL database to Amazon RDS for PostgreSQL. The company's internet link is 50 MB with a VPN in the Amazon VPC, and the Solutions Architect needs to migrate the data and synchronize the changes before the cutover. The cutover must take place within an 8-day period. What is the LEAST complex method of migrating the database securely and reliably?

- A. Order an AWS Snowball device and copy the database using the AWS Data Migration Agent
- B. When the database is available in Amazon S3, use AWS DMS to load it to Amazon RDS, and configure a job to synchronize changes before the cutover.
- C. Create an AWS DMS job to continuously replicate the data from on-premises to Amazon RDS
- D. Cutover to Amazon RDS after the data is synchronized.
- E. Order an AWS Snowball device and copy a database dump to the device
- F. After the data has been copied to Amazon S3, import it to the Amazon RDS instance
- G. Set up log shipping over a VPN to synchronize changes before the cutover.
- H. Order an AWS Snowball device and copy the database by using the AWS Schema Conversion Tool. When the data is available in Amazon S3, use AWS DMS to load it to Amazon RDS, and configure a job to synchronize changes before the cutover.

Answer: B

NEW QUESTION 94

A company's application is increasingly popular and experiencing latency because of high volume reads on the database server. The service has the following properties:

- A. A highly available REST API hosted in one region using Application Load Balancer (ALB) with auto scaling.
- A. A MySQL database hosted on an Amazon EC2 instance in a single Availability Zone.

The company wants to reduce latency, increase in-region database read performance, and have multi-region disaster recovery capabilities that can perform a live recovery automatically without any data or performance loss (HA/DR). Which deployment strategy will meet these requirements?

- A. Use AWS CloudFormation StackSets to deploy the API layer in two regions
- B. Migrate the database to an Amazon Aurora with MySQL database cluster with multiple read replicas in one region and a read replica in a different region than the source database cluster
- C. Use Amazon Route 53 health checks to trigger a DNS failover to the standby region if the health checks to the primary load balancer fail
- D. In the event of Route 53 failover, promote the cross-region database replica to be the master and build out new read replicas in the standby region.
- E. Use Amazon ElastiCache for Redis Multi-AZ with an automatic failover to cache the database read queries
- F. Use AWS OpsWorks to deploy the API layer, cache layer, and existing database layer in two regions
- G. In the event of failure, use Amazon Route 53 health checks on the database to trigger a DNS failover to the standby region if the health checks in the primary region fail
- H. Back up the MySQL database frequently, and in the event of a failure in an active region, copy the backup to the standby region and restore the standby database.
- I. Use AWS CloudFormation StackSets to deploy the API layer in two regions

- J. Add the database to an Auto Scaling group
- K. Add a read replica to the database in the second region
- L. Use Amazon Route 53 health checks on the database to trigger a DNS failover to the standby region if the health checks in the primary region fail
- M. Promote the cross-region database replica to be the master and build out new read replicas in the standby region.
- N. Use Amazon ElastiCache for Redis Multi-AZ with an automatic failover to cache the database read queries
- O. Use AWS OpsWorks to deploy the API layer, cache layer, and existing database layer in two regions
- P. Use Amazon Route 53 health checks on the ALB to trigger a DNS failover to the standby region if the health checks in the primary region fail
- Q. Back up the MySQL database frequently, and in the event of a failure in an active region, copy the backup to the standby region and restore the standby database.

Answer: A

NEW QUESTION 95

A company is adding a new approved external vendor that only supports IPv6 connectivity. The company's backend systems sit in the private subnet of an Amazon VPC. The company uses a NAT gateway to allow these systems to communicate with external vendors over IPv4. Company policy requires systems that communicate with external vendors use a security group that limits access to only approved external vendors. The virtual private cloud (VPC) uses the default network ACL.

The Systems Operator successfully assigns IPv6 addresses to each of the backend systems. The Systems Operator also updates the outbound security group to include the IPv6 CIDR of the external vendor (destination). The systems within the VPC are able to ping one another successfully over IPv6. However, these systems are unable to communicate with the external vendor.

What changes are required to enable communication with the external vendor?

- A. Create an IPv6 NAT instance
- B. Add a route for destination 0.0.0.0/0 pointing to the NAT instance.
- C. Enable IPv6 on the NAT gateway
- D. Add a route for destination ::/0 pointing to the NAT gateway.
- E. Enable IPv6 on the internet gateway
- F. Add a route for destination 0.0.0.0/0 pointing to the IGW.
- G. Create an egress-only internet gateway
- H. Add a route for destination ::/0 pointing to the gateway.

Answer: D

Explanation:

<https://docs.aws.amazon.com/vpc/latest/userguide/egress-only-internet-gateway.html>

NEW QUESTION 97

A company is using AWS to run an internet-facing production application written in Node.js. The Development team is responsible for pushing new versions of their software directly to production. The application software is updated multiple times a day. The team needs guidance from a Solutions Architect to help them deploy the software to the production fleet quickly and with the least amount of disruption to the service.

Which option meets these requirements?

- A. Prepackage the software into an AMI and then use Auto Scaling to deploy the production fleet
- B. For software changes, update the AMI and allow Auto Scaling to automatically push the new AMI to production.
- C. Use AWS CodeDeploy to push the prepackaged AMI to production
- D. For software changes, reconfigure CodeDeploy with new AMI identification to push the new AMI to the production fleet.
- E. Use AWS Elastic Beanstalk to host the production application
- F. For software changes, upload the new application version to Elastic Beanstalk to push this to the production fleet using a blue/green deployment method.
- G. Deploy the base AMI through Auto Scaling and bootstrap the software using user data
- H. For software changes, SSH to each of the instances and replace the software with the new version.

Answer: C

NEW QUESTION 98

A company has several teams, and each team has their own Amazon RDS database that totals 100 TB. The company is building a data query platform for Business Intelligence Analysts to generate a weekly business report. The new system must run ad-hoc SQL queries.

What is the MOST cost-effective solution?

- A. Create a new Amazon Redshift cluster. Create an AWS Glue ETL job to copy data from the RDS databases to the Amazon Redshift cluster. Use Amazon Redshift to run the query.
- B. Create an Amazon EMR cluster with enough core nodes. Run an Apache Spark job to copy data from the RDS databases to an Hadoop Distributed File System (HDFS). Use a local Apache Hive metastore to maintain the table definition. Use Spark SQL to run the query.
- C. Use an AWS Glue ETL job to copy all the RDS databases to a single Amazon Aurora PostgreSQL database. Run SQL queries on the Aurora PostgreSQL database.
- D. Use an AWS Glue crawler to crawl all the databases and create tables in the AWS Glue Data Catalog. Use an AWS Glue ETL Job to load data from the RDS databases to Amazon S3, and use Amazon Athena to run the queries.

Answer: C

NEW QUESTION 102

A company that provides wireless services needs a solution to store and analyze log files about user activities. Currently, log files are delivered daily to Amazon Linux on Amazon EC2 instance. A batch script is run once a day to aggregate data used for analysis by a third-party tool. The data pushed to the third-party tool is used to generate a visualization for end users. The batch script is cumbersome to maintain, and it takes several hours to deliver the ever-increasing data volumes to the third-party tool. The company wants to lower costs, and is open to considering a new tool that minimizes development effort and lowers administrative overhead. The company wants to build a more agile solution that can store and perform the analysis in near-real time, with minimal overhead. The solution needs to be cost effective and scalable to meet the company's end-user base growth.

Which solution meets the company's requirements?

- A. Develop a Python script to pull the data from Amazon EC2 in real time and store the data in Amazon S3. Use a copy command to copy data from Amazon S3

to Amazon Redshift

B. Connect a business intelligence tool running on Amazon EC2 to Amazon Redshift and create the visualizations.

C. Use an Amazon Kinesis agent running on an EC2 instance in an Auto Scaling group to collect and send the data to an Amazon Kinesis Data Firehose delivery stream

D. The Kinesis Data Firehose delivery stream will deliver the data directly to Amazon E

E. Use Kibana to visualize the data.

F. Use an in-memory caching application running on an Amazon EBS-optimized EC2 instance to capture the log data in near real-time

G. Install an Amazon ES cluster on the same EC2 instance to store the log files as they are delivered to Amazon EC2 in near real-time

H. Install a Kibana plugin to create the visualizations.

I. Use an Amazon Kinesis agent running on an EC2 instance to collect and send the data to an Amazon Kinesis Data Firehose delivery stream

J. The Kinesis Data Firehose delivery stream will deliver the data to Amazon S3. Use an AWS Lambda function to deliver the data from Amazon S3 to Amazon E

K. Use Kibana to visualize the data.

Answer: B

Explanation:

<https://docs.aws.amazon.com/firehose/latest/dev/writing-with-agents.html>

NEW QUESTION 106

A finance company is running its business-critical application on current-generation Linux EC2 instances. The application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.

Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

A. Pre-warming Elastic Load Balancers, using a bigger instance type, changing all Amazon EBS volumes to GP2 volumes.

B. Performing a one-time migration of the database cluster to Amazon RDS, and creating several additional read replicas to handle the load during end of month.

C. Using Amazon CloudWatch with AWS Lambda to change the type, size, or IOPS of Amazon EBS volumes in the cluster based on a specific CloudWatch metric.

D. Replacing all existing Amazon EBS volumes with new PIOPS volumes that have the maximum available storage size and I/O per second by taking snapshots before the end of the month and reverting back afterwards.

Answer: B

NEW QUESTION 109

A company is migrating to the cloud. It wants to evaluate the configurations of virtual machines in its existing data center environment to ensure that it can size new Amazon EC2 instances accurately. The company wants to collect metrics, such as CPU, memory, and disk utilization, and it needs an inventory of what processes are running on each instance. The company would also like to monitor network connections to map communications between servers.

Which would enable the collection of this data MOST cost effectively?

A. Use AWS Application Discovery Service and deploy the data collection agent to each virtual machine in the data center.

B. Configure the Amazon CloudWatch agent on all servers within the local environment and publish metrics to Amazon CloudWatch Logs.

C. Use AWS Application Discovery Service and enable agentless discovery in the existing virtualization environment.

D. Enable AWS Application Discovery Service in the AWS Management Console and configure the corporate firewall to allow scans over a VPN.

Answer: A

NEW QUESTION 110

A Solutions Architect must create a cost-effective backup solution for a company's 500MB source code repository of proprietary and sensitive applications. The repository runs on Linux and backs up daily to tape. Tape backups are stored for 1 year.

The current solutions are not meeting the company's needs because it is a manual process that is prone to error, expensive to maintain, and does not meet the need for a Recovery Point Objective (RPO) of 1 hour or Recovery Time Objective (RTO) of 2 hours. The new disaster recovery requirement is for backups to be stored offsite and to be able to restore a single file if needed.

Which solution meets the customer's needs for RTO, RPO, and disaster recovery with the LEAST effort and expense?

A. Replace local tapes with an AWS Storage Gateway virtual tape library to integrate with current backup software

B. Run backups nightly and store the virtual tapes on Amazon S3 standard storage in US-EAST-1. Use cross-region replication to create a second copy in US-WEST-2. Use Amazon S3 lifecycle policies to perform automatic migration to Amazon Glacier and deletion of expired backups after 1 year?

C. Configure the local source code repository to synchronize files to an AWS Storage Gateway file Amazon gateway to store backup copies in an Amazon S3 Standard bucket

D. Enable versioning on the Amazon S3 bucket

E. Create Amazon S3 lifecycle policies to automatically migrate old versions of objects to Amazon S3 Standard Infrequent Access, then Amazon Glacier, then delete backups after 1 year.

F. Replace the local source code repository storage with a Storage Gateway stored volume

G. Change the default snapshot frequency to 1 hour

H. Use Amazon S3 lifecycle policies to archive snapshots to Amazon Glacier and remove old snapshots after 1 year

I. Use cross-region replication to create a copy of the snapshots in US-WEST-2.

J. Replace the local source code repository storage with a Storage Gateway cached volume

K. Create a snapshot schedule to take hourly snapshots

L. Use an Amazon CloudWatch Events schedule expression rule to run on hourly AWS Lambda task to copy snapshots from US-EAST -1 to US-WEST-2.

Answer: B

Explanation:

<https://d1.awsstatic.com/whitepapers/aws-storage-gateway-file-gateway-for-hybrid-architectures.pdf>

NEW QUESTION 115

A company's data center is connected to the AWS Cloud over a minimally used 10-Gbps AWS Direct Connect connection with a private virtual interface to its virtual private cloud (VPC). The company internet connection is 200 Mbps and the company has a 150-TB dataset that is created each Friday. The data must be

transferred and available in Amazon S3 on Monday morning

Which is the LEAST expensive way to meet the requirements while allowing for data transfer growth?

- A. Order two 80-GB AWS Snowball appliances Offload the data to the appliances and ship them to AWS AWS will copy the data from the Snowball appliances to Amazon S3
- B. Create a VPC endpoint for Amazon S3 Copy the data to Amazon S3 by using the VPC endpoint forcing the transfer to use the Direct Connect connection
- C. Create a VPC endpoint for Amazon S3 Set up a reverse proxy farm behind a Classic Load Balancer in the VPC Copy the data to Amazon S3 using the proxy
- D. Create a public virtual interface on a Direct Connect connection and copy the data to Amazon S3 over the connection

Answer: D

NEW QUESTION 118

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 configuratio
- B. Ensure that all EC2 instances are purchased as reserved instance
- 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 EC
- E. Migrate to an Amazon Aurora MySQL Multi-AZ cluste
- F. Implement storage checks for Aurora MySQL storage utilization and an AWS Lambda function to grow the Aurora MySQL storage, as necessar
- 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 instance
- I. Ensure that the reserved instances are purchased for fixed capacity and that Auto Scaling instances run on deman
- J. Migrate to an Amazon Aurora MySQLMulti-AZ cluste
- 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 instance
- N. Ensure that Reserved instances are purchased for fixed capacity and that Auto Scaling instances run on deman
- O. Migrate to an Amazon Aurora MySQL Multi-AZ cluste
- P. Implement storage checks for Aurora MySQL storage utilization and an AWS Lambda function to grow Aurora MySQL storage, as necessar
- Q. Ensure Multi-AZ architectures are implemented.

Answer: C

NEW QUESTION 119

A Solutions Architect must establish a patching plan for a large mixed fleet of Windows and Linux servers. The patching plan must be implemented securely, be audit ready, and comply with the company's business requirements.

Which option will meet these requirements with MINIMAL effort?

- A. Install and use an OS-native patching service to manage the update frequency and release approval for all instance
- B. Use AWS Config to verify the OS state on each instance and report on any patch compliance issues.
- C. Use AWS Systems Manager on all instances to manage patchin
- D. Test patches outside of production and then deploy during a maintenance window with the appropriate approval.
- E. Use AWS OpsWorks for Chef Automate to run a set of scripts that will iterate through all instances of a given typ
- F. Issue the appropriate OS command to get and install updates on each instance, including any required restarts during the maintenance window.
- G. Migrate all applications to AWS OpsWorks and use OpsWorks automatic patching support to keep the OS up-to-date following the initial installatio
- H. Use AWS Config to provide audit and compliance reporting.

Answer: B

Explanation:

Only Systems Manager can patch both OS effectively on AWS and on premise.

NEW QUESTION 120

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 122

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 theon-premises cluste

- B. Store the data on EMRF
- 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 cluste
- F. Store the data on EMRF
- 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 theon-premises cluste
- J. Store the on EMRF
- 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 theon-premises cluste
- N. Store the data on EMRF
- 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 123

A company has an internal AWS Elastic Beanstalk worker environment inside a VPC that must access an external payment gateway API available on an HTTPS endpoint the public internet. Because of security policies, the payment gateway's Application team can grant access to only one public IP address. Which architecture will set up an Elastic Beanstalk environment to access the company's application without making multiple changes on the company's end?

- A. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a private subnet with an outbound route to a NAT gateway in a public subnet. Associate an Elastic IP address to the NAT gateway that can be whitelisted on the payment gateway application side.
- B. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a public subnet with an internet gateway. Associate an Elastic IP address to the internet gateway that can be whitelisted on the payment gateway application side.
- C. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a private subnet. Set an `https_proxy` application parameter to send outbound HTTPS connections to an EC2 proxy server deployed in a public subnet. Associate an Elastic IP address to the EC2 proxy host that can be whitelisted on the payment gateway application side.
- D. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a public subnet. Set the `https_proxy` and `no_proxy` application parameters to send non-VPC outbound HTTPS connections to an EC2 proxy server deployed in a public subnet. Associate an Elastic IP address to the EC2 proxy host that can be whitelisted on the payment gateway application side.

Answer: C

NEW QUESTION 125

A Solutions Architect is designing the storage layer for a data warehousing application. The data files are large, but they have statically placed metadata at the beginning of each file that describes the size and placement of the file's index. The data files are read in by a fleet of Amazon EC2 instances that store the index size, index location, and other category information about the data file in a database. That database is used by Amazon EMR to group files together for deeper analysis.

What would be the MOST cost-effective, high availability storage solution for this workflow?

- A. Store the data files in Amazon S3 and use Range GET for each file's metadata, then index the relevant data.
- B. Store the data files in Amazon EFS mounted by the EC2 fleet and EMR nodes.
- C. Store the data files on Amazon EBS volumes and allow the EC2 fleet and EMR to mount and unmount the volumes where they are needed.
- D. Store the content of the data files in Amazon DynamoDB tables with the metadata, index, and data as their own keys.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/API/RESTObjectGET.html>

NEW QUESTION 128

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 identity 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 131

A company has a website that enables users to upload videos. Company policy states the uploaded videos must be analyzed for restricted content. An uploaded video is placed in Amazon S3, and a message is pushed to an Amazon SQS queue with the video's location. A backend application pulls this location from Amazon SQS and analyzes the video.

The video analysis is compute-intensive and occurs sporadically during the day. The website scales with demand. The video analysis application runs on a fixed number of instances. Peak demand occurs during the holidays, so the company must add instances to the application during this time. All instances used are currently on-demand Amazon EC2 T2 instances. The company wants to reduce the cost of the current solution.

Which of the following solutions is MOST cost-effective?

- A. Keep the website on T2 instances. Determine the minimum number of website instances required during off-peak times and use Spot Instances to cover them while using Reserved Instances to cover peak demand. Use Amazon EC2 R4 and Amazon EC2 R5 Reserved Instances in an Auto Scaling group for the video analysis application.
- B. Keep the website on 12 instances. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak demand. Use Spot Fleet for the video analysis application comprised of Amazon EC2 C4 and Amazon EC2 C5 Spot Instances.
- C. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 C4 instances. Determine the minimum number of website instances required during off-peak times and use On-Demand instances to cover them while using Spot capacity to cover peak demand. Use Spot Fleet for the video analysis application comprised of C4 and Amazon EC2 C5 instances.
- D. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 R4 instances. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak demand. Use Spot Fleet for the video analysis application comprised of R4 and Amazon EC2 R5 instances.

Answer: B

NEW QUESTION 132

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 134

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 logs.
- D. Search the log data using a pre-defined set of filter patterns that machines mutating API calls.
- 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 accounts.
- H. Ensure that CloudTrail is enabled in all accounts and available AWS services.
- 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 resources.
- 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 services.
- 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 139

An organization has a write-intensive mobile application that uses Amazon API Gateway, AWS Lambda, and Amazon DynamoDB. The application has scaled well, however, costs have increased exponentially because of higher than anticipated Lambda costs. The application's use is unpredictable, but there has been a

steady 20% increase in utilization every month.

While monitoring the current Lambda functions, the Solutions Architect notices that the execution-time averages 4.5 minutes. Most of the wait time is the result of a high-latency network call to a 3-TB MySQL database server that is on-premises. A VPN is used to connect to the VPC, so the Lambda functions have been configured with a five-minute timeout.

How can the Solutions Architect reduce the cost of the current architecture?

- A. Replace the VPN with AWS Direct Connect to reduce the network latency to the on-premises MySQL database. Enable local caching in the mobile application to reduce the Lambda function invocation calls. Monitor the Lambda function performance; gradually adjust the timeout and memory properties to lower values while maintaining an acceptable execution time. Offload the frequently accessed records from DynamoDB to Amazon ElastiCache.
- B. Replace the VPN with AWS Direct Connect to reduce the network latency to the on-premises MySQL database. Cache the API Gateway results to Amazon CloudFront. Use Amazon EC2 Reserved Instances instead of Lambda. Enable Auto Scaling on EC2, and use Spot Instances during peak times. Enable DynamoDB Auto Scaling to manage target utilization.
- C. Migrate the MySQL database server into a Multi-AZ Amazon RDS for MySQL. Enable caching of the Amazon API Gateway results in Amazon CloudFront to reduce the number of Lambda function invocations. Monitor the Lambda function performance; gradually adjust the timeout and memory properties to lower values while maintaining an acceptable execution time. Enable DynamoDB Accelerator for frequently accessed records, and enable the DynamoDB Auto Scaling feature.
- D. Migrate the MySQL database server into a Multi-AZ Amazon RDS for MySQL. Enable API caching on API Gateway to reduce the number of Lambda function invocations. Continue to monitor the AWS Lambda function performance; gradually adjust the timeout and memory properties to lower values while maintaining an acceptable execution time. Enable Auto Scaling in DynamoDB.

Answer: D

NEW QUESTION 144

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 instances.
- B. Use AWS CodeDeploy to push the application code to the instances.
- C. Store the static data in Amazon EFS.
- D. Use AWS System Manager to create a new AMI with updated OS patches.
- E. Update the Auto Scaling group to use the patches AMI and replace existing unpatched instances 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 release.
- K. Use AWS CodeDeploy to push the application code to the instances.

Answer: C

NEW QUESTION 149

A company has decided to move some workloads onto AWS to create a grid environment to run market analytics. The grid will consist of many similar instances, spun-up by a job-scheduling function. Each time a large analytics workload is completed, a new VPC is deployed along with job scheduler and grid nodes. Multiple grids could be running in parallel.

Key requirements are:

- Grid instances must communicate with Amazon S3 to retrieve data to be processed.
- Grid instances must communicate with Amazon DynamoDB to track intermediate data.
- The job scheduler needs only to communicate with the Amazon EC2 API to start new grid nodes.

A key requirement is that the environment has no access to the internet, either directly or via the on-premises proxy. However, the application needs to be able to seamlessly communicate to Amazon S3, Amazon DynamoDB, and Amazon EC2 API, without the need for reconfiguration for each new deployment.

Which of the following should the Solutions Architect do to achieve this target architecture? (Choose three.)

- A. Enable VPC endpoints for Amazon S3 and DynamoDB.
- B. Disable Private DNS Name Support.
- C. Configure the application on the grid instances to use the private DNS name of the Amazon S3 endpoint.
- D. Populate the on-premises DNS server with the private IP addresses of the EC2 endpoint.
- E. Enable an interface VPC endpoint for EC2.
- F. Configure Amazon S3 endpoint policy to permit access only from the grid nodes.

Answer: ACE

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/connect-s3-vpc-endpoint/> <https://docs.aws.amazon.com/vpc/latest/userguide/vpce-interface.html>

NEW QUESTION 150

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 152

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 CM
- 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 154

A Solutions Architect is designing a multi-account structure that has 10 existing accounts. The design must meet the following requirements:

- Consolidate all accounts into one organization.
 - Allow full access to the Amazon EC2 service from the master account and the secondary accounts.
 - Minimize the effort required to add additional secondary accounts.
- Which combination of steps should be included in the solution? (Choose two.)

- A. Create an organization from the master account
- B. Send invitations to the secondary accounts from the master account
- C. Accept the invitations and create an OU.
- D. Create an organization from the master account
- E. Send a join request to the master account from each secondary account
- F. Accept the requests and create an OU.
- G. Create a VPC peering connection between the master account and the secondary account
- H. Accept the request for the VPC peering connection.
- I. Create a service control policy (SCP) that enables full EC2 access, and attach the policy to the OU.
- J. Create a full EC2 access policy and map the policy to a role in each account
- K. Trust every other account to assume the role.

Answer: AD

Explanation:

There is a concept of Permission Boundary vs Actual IAM Policies. That is, we have a concept of "Allow" vs "Grant". In terms of boundaries, we have the following three boundaries: 1. SCP 2. User/Role boundaries 3. Session boundaries (ex. AssumeRole ...) In terms of actual permission granting, we have the following: 1. Identity Policies 2. Resource Policies

NEW QUESTION 157

A company is migrating its on-premises build artifact server to an AWS solution. The current system consists of an Apache HTTP server that serves artifacts to clients on the local network, restricted by the perimeter firewall. The artifact consumers are largely build automation scripts that download artifacts via anonymous HTTP, which the company will be unable to modify within its migration timetable.

The company decides to move the solution to Amazon S3 static website hosting. The artifact consumers will be migrated to Amazon EC2 instances located within both public and private subnets in a virtual private cloud (VPC).

Which solution will permit the artifact consumers to download artifacts without modifying the existing automation scripts?

- A. Create a NAT gateway within a public subnet of the VPC
- B. Add a default route pointing to the NAT gateway into the route table associated with the subnets containing consumer
- C. Configure the bucket policy to allow the s3:ListBucket and s3:GetObject actions using the condition IpAddress and the condition key aws:SourceIp matching the elastic IP address of the NAT gateway.
- D. Create a VPC endpoint and add it to the route table associated with subnets containing consumers. Configure the bucket policy to allow s3:ListBucket and s3:GetObject actions using the condition StringEquals and the condition key aws:sourceVpce matching the identification of the VPC endpoint.
- E. Create an IAM role and instance profile for Amazon EC2 and attach it to the instances that consume build artifact
- F. Configure the bucket policy to allow the s3:ListBucket and s3:GetObjects actions for the principal matching the IAM role created.
- G. Create a VPC endpoint and add it to the route table associated with subnets containing consumers. Configure the bucket policy to allow s3:ListBucket and s3:GetObject actions using the condition IpAddress and the condition key aws:SourceIp matching the VPC CIDR block.

Answer: B

NEW QUESTION 159

A financial services company is moving to AWS and wants to enable Developers to experiment and innovate while preventing access to production applications

The company has the following requirements

- Production workloads cannot be directly connected to the internet
- All workloads must be restricted to the us-west-2 and eu-central-1 Regions
- Notification should be sent when Developer sandboxes exceed \$500 in AWS spending monthly

Which combination of actions needs to be taken to create a multi-account structure that meets the company's requirements'? (Select THREE)

- A. Create accounts for each production workload within an organization in AWS Organizations Place the production accounts within an organizational unit (OU) For each account delete the default VPC Create an SCP with a Deny rule for the attach an internet gateway and create a default VPC actions Attach the SCP to the OU for the production accounts
- B. Create accounts for each production workload within an organization in AWS Organizations Place the production accounts within an organizational unit (OU) Create an SCP with a Deny rule on the attach an internet gateway action Create an SCP with a Deny rule to prevent use of the default VPC Attach the SCPs to the OU for the production accounts
- C. Create a SCP containing a Deny Effect for cloudfront". lam:*, route53* and support* with a StringNotEquals condition on an aws RequestedRegion condition key with us-west-2 and eu-central-1 values Attach the SCP to the organization's root.
- D. Create an IAM permission boundary containing a Deny Effect for cloudfront'. lam * route53' and support" with a StringNotEquals condition on an aws RequestedRegion condition key with us-west 2 and eu-central-1 values Attach the permission boundary to an IAM group containing the development and production users.
- E. Create accounts for each development workload within an organization m AWS Organizations Place the development accounts within an organizational unit (OU) Create a custom AWS Config rule to deactivate all (AM users when an account's monthly bill exceeds \$500.
- F. Create accounts for each development workload within an organization in AWS Organizations Place the development accounts within an organizational unit (OU) Create a budget within AWS Budgets for each development account to monitor and report on monthly spending exceeding \$500.

Answer: ABD

NEW QUESTION 162

A utility company wants to collect usage data every 5 minutes from its smart meters to facilitate time-of-use metering When a meter sends data to AWS the data is sent to Amazon API Gateway, processed by an AWS Lambda function and stored in an Amazon DynamoDB table During the pilot phase, the Lambda functions took from 3 to 5 seconds to complete

As more smart meters are deployed, the Engineers notice the Lambda functions are taking from 1 to 2 minutes to complete The functions are also increasing in duration as new types of metrics are collected from the devices There are many ProvisionedThroughputExceededException errors while performing PUT operations on DynamoDB and there are also many TooManyRequests Exception errors from Lambda.

Which combination of changes will resolve these issues? (Select TWO)

- A. increase the write capacity units to the DynamoDB table
- B. Increase the memory available to the Lambda functions
- C. Increase the payload size from the smart meters to send more data
- D. Stream the data into an Amazon Kinesis data stream from API Gateway and process the data in batches
- E. Collect data in an Amazon SQS FIFO queue, which triggers a Lambda function to process each message

Answer: AB

NEW QUESTION 165

A company is running a web application with On-Demand Amazon EC2 instances in Auto Scaling groups that scale dynamically based on custom metrics After extensive testing the company determines that the m5 2xlarge instance size is optimal for the workload Application data is stored in db r4 4xlarge Amazon RDS instances that are confirmed to be optimal The traffic to the web application spikes randomly during the day

What other cost-optimization methods should the company implement to further reduce costs without impacting the reliability of the application?

- A. Double the instance count in the Auto Scaling groups and reduce the instance size to m5 large
- B. Reserve capacity for the RDS database and the minimum number of EC2 instances that are constantly running
- C. Reduce the RDS instance size to db r4 xlarge and add five equivalents sized read replicas to provide reliability
- D. Reserve capacity for all EC2 instances and leverage Spot Instance pricing for the RDS database

Answer: B

NEW QUESTION 169

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 identity 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 172

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 174

A media storage application uploads user photos to Amazon S3 for processing. End users are reporting that some uploaded photos are not being processed properly. The Application Developers trace the logs and find that AWS Lambda is experiencing execution issues when thousands of users are on the system simultaneously. Issues are caused by:

Limits around concurrent executions.

The performance of Amazon DynamoDB when saving data.

Which actions can be taken to increase the performance and reliability of the application? (Choose two.)

- A. Evaluate and adjust the read capacity units (RCUs) for the DynamoDB tables.
- B. Evaluate and adjust the write capacity units (WCUs) for the DynamoDB tables.
- C. Add an Amazon ElastiCache layer to increase the performance of Lambda functions.
- D. Configure a dead letter queue that will reprocess failed or timed-out Lambda functions.
- E. Use S3 Transfer Acceleration to provide lower-latency access to end users.

Answer: BD

Explanation:

B:
<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.h>
D: <https://aws.amazon.com/blogs/compute/robust-serverless-application-design-with-aws-lambda-dlq/c>

NEW QUESTION 175

A company is planning the migration of several lab environments used for software testing. An assortment of custom tooling is used to manage the test runs for each lab. The labs use immutable infrastructure for the software test runs, and the results are stored in a highly available SQL database cluster. Although completely rewriting the custom tooling is out of scope for the migration project, the company would like to optimize workloads during the migration. Which application migration strategy meets this requirement?

- A. Re-host
- B. Re-platform
- C. Re-factor/re-architect
- D. Retire

Answer: B

Explanation:

<https://aws.amazon.com/blogs/enterprise-strategy/6-strategies-for-migrating-applications-to-the-cloud/>

NEW QUESTION 179

A company has an application that runs a web service on Amazon EC2 instances and stores .jpg images in Amazon S3. The web traffic has a predictable baseline, but often demand spikes unpredictably for short periods of time. The application is loosely coupled and stateless. The .jpg images stored in Amazon S3 are accessed frequently for the first 15 to 20 days, they are seldom accessed thereafter but always need to be immediately available. The CIO has asked to find ways to reduce costs.

Which of the following options will reduce costs? (Choose two.)

- A. Purchase Reserved instances for baseline capacity requirements and use On-Demand instances for the demand spikes.
- B. Configure a lifecycle policy to move the .jpg images on Amazon S3 to S3 IA after 30 days.
- C. Use On-Demand instances for baseline capacity requirements and use Spot Fleet instances for the demand spikes.
- D. Configure a lifecycle policy to move the .jpg images on Amazon S3 to Amazon Glacier after 30 days.
- E. Create a script that checks the load on all web servers and terminates unnecessary On-Demand instances.

Answer: AB

NEW QUESTION 181

A public retail web application uses an Application Load Balancer (ALB) in front of Amazon EC2 instances running across multiple Availability Zones (AZs) in a Region backed by an Amazon RDS MySQL Multi-AZ deployment. Target group health checks are configured to use HTTP and pointed at the product catalog page. Auto Scaling is configured to maintain the web fleet size based on the ALB health check. Recently, the application experienced an outage. Auto Scaling continuously replaced the instances during the outage. A subsequent investigation determined that the web server metrics were within the normal range, but the database tier was experiencing high load, resulting in severely elevated query response times. Which of the following changes together would remediate these issues while improving monitoring capabilities for the availability and functionality of the entire application stack for future growth? (Select TWO.)

- A. Configure read replicas for Amazon RDS MySQL and use the single reader endpoint in the web application to reduce the load on the backend database tier.
- B. Configure the target group health check to point at a simple HTML page instead of a product catalog page and the Amazon Route 53 health check against the product page to evaluate full application functionalit
- C. Configure Amazon CloudWatch alarms to notify administrators when the site fails.
- D. Configure the target group health check to use a TCP check of the Amazon EC2 web server and the Amazon Route 53 health check against the product page to evaluate full application functionalit
- E. Configure Amazon CloudWatch alarms to notify administrators when the site fails.
- F. Configure an Amazon CloudWatch alarm for Amazon RDS with an action to recover a high-load, impaired RDS instance in the database tier.
- G. Configure an Amazon ElastiCache cluster and place it between the web application and RDS MySQL instances to reduce the load on the backend database tier.

Answer: BE

NEW QUESTION 183

An auction website enables users to bid on collectible items. The auction rules require that each bid is processed only once and in the order it was received. The current implementation is based on a fleet of Amazon EC2 web servers that write bid records into Amazon Kinesis Data Streams. A single t2.large instance has a cron job that runs the bid processor, which reads incoming bids from Kinesis Data Streams and processes each bid. The auction site is growing in popularity, but users are complaining that some bids are not registering.

Troubleshooting indicates that the bid processor is too slow during peak demand hours, sometimes crashes while processing, and occasionally loses track of which records is being processed.

What changes should make the bid processing more reliable?

- A. Refactor the web application to use the Amazon Kinesis Producer Library (KPL) when posting bids to Kinesis Data Stream
- B. Refactor the bid processor to flag each record in Kinesis Data Streams as being unread, processing, and processe
- C. At the start of each bid processing run, scan Kinesis Data Streams for unprocessed records.
- D. Refactor the web application to post each incoming bid to an Amazon SNS topic in place of Kinesis Data Stream
- E. Configure the SNS topic to trigger an AWS Lambda function that processes each bid as soon as a user submits it.
- F. Refactor the web application to post each incoming bid to an Amazon SQS FIFO queue in place of Kinesis Data Stream
- G. Refactor the bid processor to continuously the SQS queu
- H. Place the bid processing EC2 instance in an Auto Scaling group with a minimum and a maximum size of 1.
- I. Switch the EC2 instance type from t2.large to a larger general compute instance typ
- J. Put the bid processor EC2 instances in an Auto Scaling group that scales out the number of EC2 instances running the bid processor, based on the IncomingRecords metric in Kinesis Data Streams.

Answer: C

Explanation:

FIFO is better in this case compared to Kinesis, as it guarantee the order of the bid. Min Max 1, is okay as the SQS will hold the queue in case of failure of the instance, till it come back again.

NEW QUESTION 185

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 dat
- 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 187

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

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

100% Actual & Verified — Instant Download, Please Click
[Order The SAP-C01 Practice Test Here](#)