

## Exam Questions SAP-C01

AWS Certified Solutions Architect- Professional

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



#### NEW QUESTION 1

A company is currently running a production workload on AWS that is very I/O intensive. Its workload consists of a single tier with 10 c4.8xlarge instances, each with 2 TB gp2 volumes. The number of processing jobs has recently increased, and latency has increased as well. The team realizes that they are constrained on the IOPS. For the application to perform efficiently, they need to increase the IOPS by 3,000 for each of the instances. Which of the following designs will meet the performance goal MOST cost effectively?

- A. Change the type of Amazon EBS volume from gp2 to io1 and set provisioned IOPS to 9,000.
- B. Increase the size of the gp2 volumes in each instance to 3 TB.
- C. Create a new Amazon EFS file system and move all the data to this new file system
- D. Mount this file system to all 10 instances.
- E. Create a new Amazon S3 bucket and move all the data to this new bucket
- F. Allow each instance to access this S3 bucket and use it for storage.

**Answer: B**

#### NEW QUESTION 2

A Solutions Architect is building a containerized .NET Core application that will run in AWS Fargate. The backend of the application requires Microsoft SQL Server with high availability. All tiers of the application must be highly available. The credentials used for the connection string to SQL Server should not be stored on disk within the .NET Core front-end containers.

Which strategies should the Solutions Architect use to meet these requirements?

- A. Set up SQL Server to run in Fargate with Service Auto Scaling. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to SQL Server running in Fargate. Specify the ARN of the secret in AWS Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection string. Set up the .NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.
- B. Create a Multi-AZ deployment of SQL Server on Amazon RDS. Create a secret in AWS Secrets Manager for the credentials to the RDS database. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to the RDS database in Secrets Manager. Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection string. Set up the .NET Core service in Fargate using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.
- C. Create an Auto Scaling group to run SQL Server on Amazon EC2. Create a secret in AWS Secrets Manager for the credentials to SQL Server running on EC2. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to SQL Server on EC2. Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection string. Set up the .NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.
- D. Create a Multi-AZ deployment of SQL Server on Amazon RDS. Create a secret in AWS Secrets Manager for the credentials to the RDS database. Create non-persistent empty storage for the .NET Core containers in the Fargate task definition to store the sensitive information. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to the RDS database in Secrets Manager. Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be written to the non-persistent empty storage on startup for reading into the application to construct the connection.

**Answer: C**

#### NEW QUESTION 3

A retail company processes point-of-sale data on application servers in its data center and writes outputs to Amazon DynamoDB table. The data center is connected to the company's VPC with an AWS Direct Connect (DX) connection, and the application servers require a consistent network connection at speed greater than 2 Gbps.

The company decides that the DynamoDB table needs to be highly available and fault tolerant. The company policy states that the data should be available across two regions.

What changes should the company make to meet these requirements?

- A. Establish a second DX connection for redundancy.
- B. Use DynamoDB global tables to replicate data to a second Region. Modify the application to fail over to the second Region.
- C. Use an AWS managed VPN as a backup to DX.
- D. Create an identical DynamoDB table in a second Region.
- E. Modify the application to replicate data to both regions.
- F. Establish a second DX connection for redundancy.
- G. Create an identical DynamoDB table in a second Region.
- H. Enable DynamoDB auto scaling to manage throughput capacity.
- I. Modify the application to write to the second Region.
- J. Use AWS managed VPN as a backup to DX.
- K. Create an identical DynamoDB table in a second Region. Enable DynamoDB streams to capture changes to the table.
- L. Use AWS Lambda to replicate changes to the second Region.

**Answer: A**

#### NEW QUESTION 4

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 5**

A company wants to follow its website on AWS using serverless architecture design patterns for global customers. The company has outlined its requirements as follow:

- The website should be responsive.
- The website should offer minimal latency.
- The website should be highly available.
- Users should be able to authenticate through social identity providers such as Google, Facebook, and Amazon.
- There should be baseline DDoS protections for spikes in traffic.

How can the design requirements be met?

- A. Use Amazon CloudFront with Amazon ECS for hosting the website
- B. Use AWS Secrets Manager for provide user management and authentication function
- C. Use ECS Docker containers to build an API.
- D. Use Amazon Route 53 latency routing with an Application Load Balancer and AWS Fargate in different regions for hosting the website
- E. use Amazon Cognito to provide user management and authentication function
- F. Use Amazon EKS containers.
- G. Use Amazon CloudFront with Amazon S3 for hosting static web resource
- H. Use Amazon Cognito to provide user management authentication function
- I. Use Amazon API Gateway with AWS Lambda to build an API.
- J. Use AWS Direct Connect with Amazon CloudFront and Amazon S3 for hosting static web resource. Use Amazon Cognito to provide user management authentication function
- K. Use AWS Lambda to build an API.

**Answer:** C

**NEW QUESTION 6**

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

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

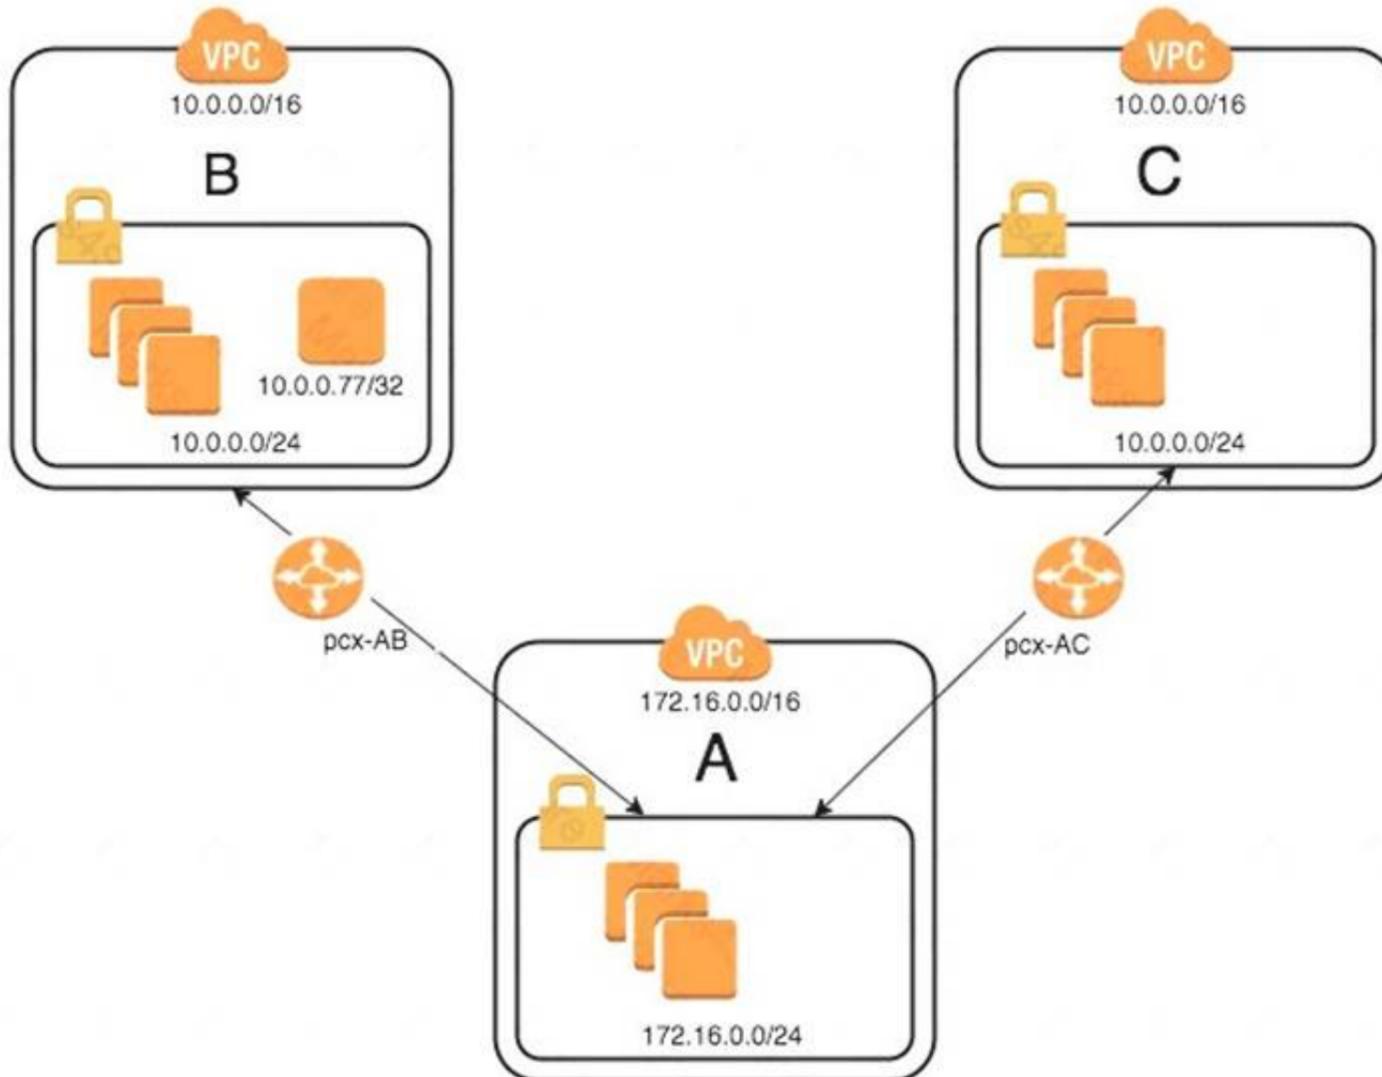
**Answer:** A

**Explanation:**

[https://docs.aws.amazon.com/AmazonECS/latest/developerguide/scheduling\\_tasks.html](https://docs.aws.amazon.com/AmazonECS/latest/developerguide/scheduling_tasks.html)

**NEW QUESTION 7**

An organization has recently grown through acquisitions. Two of the purchased companies use the same IP CIDR range. There is a new short-term requirement to allow AnyCompany A (VPC-A) to communicate with a server that has the IP address 10.0.0.77 in AnyCompany B (VPC-B). AnyCompany A must also communicate with all resources in AnyCompany C (VPC-C). The Network team has created the VPC peer links, but it is having issues with communications between VPC-A and VPC-B. After an investigation, the team believes that the routing tables in the VPCs are incorrect.



What configuration will allow AnyCompany A to communicate with AnyCompany C in addition to the database in AnyCompany B?

- A. On VPC-A, create a static route for the VPC-B CIDR range (10.0.0.0/24) across VPC peer pcx-AB. Create a static route of 10.0.0.0/16 across VPC peer pcx-AC. On VPC-B, create a static route for VPC-A CIDR (172.16.0.0/24) on peer pcx-AB. On VPC-C, create a static route for VPC-A CIDR (172.16.0.0/24) across peer pcx-AC.
- B. On VPC-A, enable dynamic route propagation on pcx-AB and pcx-AC. On VPC-B, enable dynamic route propagation and use security groups to allow only the IP address 10.0.0.77/32 on VPC peer pcx-AB. On VPC-C, enable dynamic route propagation with VPC-A on peer pcx-AC.
- C. On VPC-A, create network access control lists that block the IP address 10.0.0.77/32 on VPC peer pcx-AC. On VPC-A, create a static route for VPC-B CIDR (10.0.0.0/24) on pcx-AB and a static route for VPC-C CIDR (10.0.0.0/24) on pcx-AC. On VPC-B, create a static route for VPC-A CIDR (172.16.0.0/24) across peer pcx-AB. On VPC-C, create a static route for VPC-A CIDR (172.16.0.0/24) across peer pcx-AC.
- D. On VPC-A, create a static route for the VPC-B CIDR (10.0.0.77/32) database across VPC peer pcx-AB. Create a static route for the VPC-C CIDR on VPC peer pcx-AC. On VPC-B, create a static route for VPC-A CIDR (172.16.0.0/24) on peer pcx-AB. On VPC-C, create a static route for VPC-A CIDR (172.16.0.0/24) across peer pcx-AC.

Answer: D

**NEW QUESTION 8**

A company runs a memory-intensive analytics application using on-demand Amazon EC2 compute optimized instance. The application is used continuously and application demand doubles during working hours. The application currently scales based on CPU usage. When scaling in occurs, a lifecycle hook is used because the instance requires 4 minutes to clean the application state before terminating.

Because users reported poor performance during working hours, scheduled scaling actions were implemented so additional instances would be added during working hours. The Solutions Architect has been asked to reduce the cost of the application. Which solution is MOST cost-effective?

- A. Use the existing launch configuration that uses C5 instances, and update the application AMI to include the Amazon CloudWatch agent
- B. Change the Auto Scaling policies to scale based on memory utilization
- C. Use Reserved Instances for the number of instances required after working hours, and use Spot Instances to cover the increased demand during working hours.
- D. Update the existing launch configuration to use R5 instances, and update the application AMI to include SSM Agent
- E. Change the Auto Scaling policies to scale based on memory utilization
- F. Use Reserved instances for the number of instances required after working hours, and use Spot Instances with On-Demand instances to cover the increased demand during working hours.
- G. Use the existing launch configuration that uses C5 instances, and update the application AMI to include SSM Agent
- H. Leave the Auto Scaling policies to scale based on CPU utilization
- I. Use scheduled Reserved Instances for the number of instances required after working hours, and use Spot Instances to cover the increased demand during work hours.
- J. Create a new launch configuration using R5 instances, and update the application AMI to include the Amazon CloudWatch agent
- K. Change the Auto Scaling policies to scale based on memory utilization
- L. Use Reserved Instances for the number of instances required after working hours, and use Standard Reserved Instances with On-Demand Instances to cover the increased demand during working hours.

Answer: D

**Explanation:**

[https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring\\_ec2.html](https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring_ec2.html)

### NEW QUESTION 9

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

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

**Answer: D**

### NEW QUESTION 10

While debugging a backend application for an IoT system that supports globally distributed devices a Solutions Architect notices that stale data is occasionally being sent to user devices. Devices often share data, and stale data does not cause issues in most cases. However, device operations are disrupted when a device reads the stale data after an update. The global system has multiple identical application stacks deployed in different AWS Regions. If a user device travels out of its home geographic region it will always connect to the geographically closest AWS Region to write or read data. The same data is available in all supported AWS Regions using an Amazon DynamoDB global table. What change should be made to avoid causing disruptions in device operations'?

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

**Answer: B**

### NEW QUESTION 10

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 12

A company has a serverless application comprised of Amazon CloudFront, Amazon API Gateway, and AWS Lambda functions. The current deployment process of the application code is to create a new version number of the Lambda function and run an AWS CLI script to update. If the new function version has errors, another CLI script reverts by deploying the previous working version of the function. The company would like to decrease the time to deploy new versions of the application logic provided by the Lambda functions, and also reduce the time to detect and revert when errors are identified. How can this be accomplished?

- A. Create and deploy nested AWS CloudFormation stacks with the parent stack consisting of the AWS CloudFront distribution and API Gateway, and the child stack containing the Lambda function
- B. For changes to Lambda, create an AWS CloudFormation change set and deploy; if errors are triggered, revert the AWS CloudFormation change set to the previous version.
- C. Use AWS SAM and built-in AWS CodeDeploy to deploy the new Lambda version, gradually shift traffic to the new version, and use pre-traffic and post-traffic test functions to verify code
- D. Rollback if Amazon CloudWatch alarms are triggered.
- E. Refactor the AWS CLI scripts into a single script that deploys the new Lambda version
- F. When deployment is completed, the script tests execution

- G. If errors are detected, revert to the previous Lambda version.
- H. Create and deploy an AWS CloudFormation stack that consists of a new API Gateway endpoint that references the new Lambda version.
- I. Change the CloudFront origin to the new API Gateway endpoint, monitor errors and if detected, change the AWS CloudFront origin to the previous API Gateway endpoint.

**Answer: B**

**Explanation:**

<https://aws.amazon.com/about-aws/whats-new/2017/11/aws-lambda-supports-traffic-shifting-and-phased-deploy> <https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverless>

**NEW QUESTION 13**

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 SSO (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 15**

A company is migrating its marketing website and content management system from an on-premises data center to AWS. The company wants the AWS application to be developed in a VPC with Amazon EC2 instances used for the web servers and an Amazon RDS instance for the database. The company has a runbook document that describes the installation process of the on-premises system. The company would like to base the AWS system on the processes referenced in the runbook document. The runbook document describes the installation and configuration of the operating systems, network settings, the website, and content management system software on the servers. After the migration is complete, the company wants to be able to make changes quickly to take advantage of other AWS features. How can the application and environment be deployed and automated in AWS, while allowing for future changes?

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

**Answer: D**

**NEW QUESTION 16**

A company runs its containerized batch jobs on Amazon ECS. The jobs are scheduled by submitting a container image, a task definition, and the relevant data to an Amazon S3 bucket. Container images may be unique per job. Running the jobs as quickly as possible is of utmost importance, so submitting jobs artifacts to the S3 bucket triggers the job to run immediately. Sometimes there may be no jobs running at all. However, jobs of any size can be submitted with no prior warning to the IT Operations team. Job definitions include CPU and memory resource requirements. What solution will allow the batch jobs to complete as quickly as possible after being scheduled?

- A. Schedule the jobs on an Amazon ECS cluster using the Amazon EC2 launch type.
- B. Use Service Auto Scaling to increase or decrease the number of running tasks to suit the number of running jobs.
- C. Schedule the jobs directly on EC2 instances.
- D. Use Reserved Instances for the baseline minimum load, and use On-Demand Instances in an Auto Scaling group to scale up the platform based on demand.
- E. Schedule the jobs on an Amazon ECS cluster using the Fargate launch type.
- F. Use Service Auto Scaling to increase or decrease the number of running tasks to suit the number of running jobs.
- G. Schedule the jobs on an Amazon ECS cluster using the Fargate launch type.
- H. Use Spot Instances in an Auto Scaling group to scale the platform based on demand.
- I. Use Service Auto Scaling to increase or decrease the number of running tasks to suit the number of running jobs.

**Answer: C**

**NEW QUESTION 21**

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 23

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

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

**Answer: A**

#### Explanation:

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

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

#### NEW QUESTION 24

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

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

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

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

**Answer: C**

#### NEW QUESTION 25

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 queu
- 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 26

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 rightly using AWS Systems Manager Run Command to search (or credentials on thedevelopment 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 28

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 30

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 andpresentation 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 35

A company has an existing on-premises three-tier web application. The Linux web servers serve content from a centralized file share on a NAS server because the content is refreshed several times a day from various sources. The existing infrastructure is not optimized and the company would like to move to AWS in order to gain the ability to scale resources up and down in response to load. On-premises and AWS resources are connected using AWS Direct Connect. How can the company migrate the web infrastructure to AWS without delaying the content refresh process?

- A. Create a cluster of web server Amazon EC2 instances behind a Classic Load Balancer on AW
- B. Share an Amazon EBS volume among all instances for the conten
- C. Schedule a periodic synchronization of this volume and the NAS server.
- D. Create an on-premises file gateway using AWS Storage Gateway to replace the NAS server and replicate content to AW

- E. On the AWS side, mount the same Storage Gateway bucket to each web server Amazon EC2 instance to serve the content.
- F. Expose an Amazon EFS share to on-premises users to serve as the NAS serv
- G. Mount the same EFS share to the web server Amazon EC2 instances to serve the content.
- H. Create web server Amazon EC2 instances on AWS in an Auto Scaling grou
- I. Configure a nightly process where the web server instances are updated from the NAS server.

**Answer: C**

**Explanation:**

File gateway is limited by performance its gateway instance, whether EC2 or On-premises, Cache will get filled up fast if not properly configured, For large number of EC2 instances EFS scales better. So, bottom line is File Storage gateway is for legacy applications and you have to add cost of large gateway instances before comparing it to same quantity of EFS storage. [https://www.reddit.com/r/aws/comments/82pyop/storage\\_gateway\\_vs\\_efs/](https://www.reddit.com/r/aws/comments/82pyop/storage_gateway_vs_efs/)  
<https://docs.aws.amazon.com/efs/latest/ug/efs-onpremises.html>

**NEW QUESTION 37**

A Solutions Architect wants to make sure that only AWS users or roles with suitable permissions can access a new Amazon API Gateway endpoint The Solutions Architect wants an end-to-end view of each request to analyze the latency of the request and create service maps How can the Solutions Architect design the API Gateway access control and perform request inspections?

- A. For the API Gateway method set the authorization to AWSIAM Then, give the IAM user or role execute-api Invoke permission on the REST API resource Enable the API caller to sign requests with AWS Signature when accessing the endpoint Use AWS X-Ray to trace and analyze user requests to API Gateway
- B. For the API Gateway resource set CORS to enabled and only return the company's domain mAccess-Control-Allow-Origin headers Then give the IAM user or role execute-api Invoke permission on the REST API resource Use Amazon CloudWatch to trace and analyze user requests to API Gateway
- C. Create an AWS Lambda function as the custom authorizer ask the API client to pass the key and secret when making the call and then use Lambda to validate the key/secret pair against the IAM system Use AWS X-Ray to trace and analyze user requests to API Gateway
- D. Create a client certificate for API Gateway Distribute the certificate to the AWS users and roles that need to access the endpoint Enable the API caller to pass the client certificate when accessing the endpoint Use Amazon CloudWatch to trace and analyze user requests to API Gateway.

**Answer: D**

**NEW QUESTION 40**

A Development team is deploying new APIs as serverless applications within a company. The team is currently using the AWS Management Console to provision Amazon API Gateway, AWS Lambda, and Amazon DynamoDB resources. A Solutions Architect has been tasked with automating the future deployments of these serverless APIs.

How can this be accomplished?

- A. Use AWS CloudFormation with a Lambda-backed custom resource to provision API Gatewa
- B. Use the AWS::DynamoDB::Table and AWS::Lambda::Function resources to create the Amazon DynamoDB table and Lambda function
- C. Write a script to automate the deployment of the CloudFormation template.
- D. Use the AWS Serverless Application Model to define the resource
- E. Upload a YAML template and application files to the code repositor
- F. Use AWS CodePipeline to connect to the code repository and to create an action to build using AWS CodeBuil
- G. Use the AWS CloudFormation deployment provider in CodePipeline to deploy the solution.
- H. Use AWS CloudFormation to define the serverless applicatio
- I. Implement versioning on the Lambda functions and create aliases to point to the version
- J. When deploying, configure weights to implement shifting traffic to the newest version, and gradually update the weights as traffic moves over.
- K. Commit the application code to the AWS CodeCommit code repositor
- L. Use AWS CodePipeline and connect to the CodeCommit code repositor
- M. Use AWS CodeBuild to build and deploy the Lambda functions using AWS CodeDeplo
- N. Specify the deployment preference type in CodeDeploy to gradually shift traffic over to the new version.

**Answer: B**

**Explanation:**

<https://aws-quickstart.s3.amazonaws.com/quickstart-trek10-serverless-enterprise-cicd/doc/serverless-cicd-for-th>  
<https://aws.amazon.com/quickstart/architecture/serverless-cicd-for-enterprise/>

**NEW QUESTION 41**

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 AWSTrusted Advisor checks to continue monitoring the environment for future savings.
- H. Move the entire website to Amazon S3 using the S3 website hosting featur
- 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 applicatio
- K. Move all images and video files to Amazon EF
- L. Create an Amazon CloudFront distribution that points to the EFS shar
- M. Reserve the m4.4xl instances needed to meet base performance requirements.

**Answer: B**

#### NEW QUESTION 44

A three-tier web application runs on Amazon EC2 instances. Cron daemons are used to trigger scripts that collect the web server, application, and database logs and send them to a centralized location every hour. Occasionally, scaling events or unplanned outages have caused the instances to stop before the latest logs were collected, and the log files were lost.

Which of the following options is the MOST reliable way of collecting and preserving the log files?

- A. Update the cron jobs to run every 5 minutes instead of every hour to reduce the possibility of log messages being lost in an outage.
- B. Use Amazon CloudWatch Events to trigger Amazon Systems Manager Run Command to invoke the log collection scripts more frequently to reduce the possibility of log messages being lost in an outage.
- C. Use the Amazon CloudWatch Logs agent to stream log messages directly to CloudWatch Logs. Configure the agent with a batch count of 1 to reduce the possibility of log messages being lost in an outage.
- D. Use Amazon CloudWatch Events to trigger AWS Lambda to SSH into each running instance and invoke the log collection scripts more frequently to reduce the possibility of log messages being lost in an outage.

**Answer:** C

#### Explanation:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/AgentReference.html>

#### NEW QUESTION 46

A company is implementing a multi-account strategy; however, the Management team has expressed concerns that services like DNS may become overly complex. The company needs a solution that allows private DNS to be shared among virtual private clouds (VPCs) in different accounts. The company will have approximately 50 accounts in total.

What solution would create the LEAST complex DNS architecture and ensure that each VPC can resolve all AWS resources?

- A. Create a shared services VPC in a central account, and create a VPC peering connection from the shared services VPC to each of the VPCs in the other account
- B. Within Amazon Route 53, create a privately hosted zone in the shared services VPC and resource record sets for the domain and subdomains. Programmatically associate other VPCs with the hosted zone.
- C. Create a VPC peering connection among the VPCs in all account
- D. Set the VPC attributes enableDnsHostnames and enableDnsSupport to "true" for each VP
- E. Create an Amazon Route 53 private zone for each VP
- F. Create resource record sets for the domain and subdomain
- G. Programmatically associate the hosted zones in each VPC with the other VPCs.
- H. Create a shared services VPC in a central account
- I. Create a VPC peering connection from the VPCs in other accounts to the shared services VP
- J. Create an Amazon Route 53 privately hosted zone in the shared services VPC with resource record sets for the domain and subdomain
- K. Allow UDP and TCP port 53 over the VPC peering connections.
- L. Set the VPC attributes enableDnsHostnames and enableDnsSupport to "false" in every VP
- M. Create an AWS Direct Connect connection with a private virtual interface
- N. Allow UDP and TCP port 53 over the virtual interface
- O. Use the on-premises DNS servers to resolve the IP addresses in each VPC on AWS.

**Answer:** A

#### Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/centralized-dns-management-of-hybrid-cloud-w>

#### NEW QUESTION 49

A company has multiple AWS accounts hosting IT applications. An Amazon CloudWatch Logs agent is installed on all Amazon EC2 instances. The company wants to aggregate all security events in a centralized AWS account dedicated to log storage.

Security Administrators need to perform near-real-time gathering and correlating of events across multiple AWS accounts.

Which solution satisfies these requirements?

- A. Create a Log Audit IAM role in each application AWS account with permissions to view CloudWatch Logs, configure an AWS Lambda function to assume the Log Audit role, and perform an hourly export of CloudWatch Logs data to an Amazon S3 bucket in the logging AWS account.
- B. Configure CloudWatch Logs streams in each application AWS account to forward events to CloudWatch Logs in the logging AWS account
- C. In the logging AWS account, subscribe an Amazon Kinesis Data Firehose stream to Amazon CloudWatch Events, and use the stream to persist log data in Amazon S3.
- D. Create Amazon Kinesis Data Streams in the logging account, subscribe the stream to CloudWatch Logs streams in each application AWS account, configure an Amazon Kinesis Data Firehose delivery stream with the Data Streams as its source, and persist the log data in an Amazon S3 bucket inside the logging AWS account.
- E. Configure CloudWatch Logs agents to publish data to an Amazon Kinesis Data Firehose stream in the logging AWS account, use an AWS Lambda function to read messages from the stream and push messages to Data Firehose, and persist the data in Amazon S3.

**Answer:** C

#### Explanation:

The solution uses Amazon Kinesis Data Streams and a log destination to set up an endpoint in the logging account to receive streamed logs and uses Amazon Kinesis Data Firehose to deliver log data to the Amazon Simple Storage Solution (S3) bucket. Application accounts will subscribe to stream all (or part) of their Amazon CloudWatch logs to a defined destination in the logging account via subscription filters. <https://aws.amazon.com/blogs/architecture/central-logging-in-multi-account-environments/>

#### NEW QUESTION 54

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 Analytic
- 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 55**

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

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 56**

A company is using AWS for production and development workloads. Each business unit has its own AWS account for production, and a separate AWS account to develop and deploy its applications. The Information Security department has introduced new security policies that limit access for terminating certain Amazon EC2 instances in all accounts to a small group of individuals from the Security team.

How can the Solutions Architect meet these requirements?

- A. Create a new IAM policy that allows access to those EC2 instances only for the Security team
- B. Apply this policy to the AWS Organizations master account.
- C. Create a new tag-based IAM policy that allows access to these EC2 instances only for the Security team. Tag the instances appropriately, and apply this policy in each account.
- D. Create an organizational unit under AWS Organization
- E. Move all the accounts into this organizational unit and use SCP to apply a whitelist policy to allow access to these EC2 instances for the Security team only.
- F. Set up SAML federation for all accounts in AWS
- G. Configure SAML so that it checks for the service API call before authenticating the user
- H. Block SAML from authenticating API calls if anyone other than the Security team accesses these instances.

**Answer:** B

**NEW QUESTION 59**

A company runs a dynamic mission-critical web application that has an SLA of 99.99%. Global application users access the application 24/7. The application is currently hosted on premises and routinely fails to meet its SLA, especially when millions of users access the application concurrently. Remote users complain of latency.

How should this application be redesigned to be scalable and allow for automatic failover at the lowest cost?

- A. Use Amazon Route 53 failover routing with geolocation-based routing
- B. Host the website on automatically scaled Amazon EC2 instances behind an Application Load Balancer with an additional Application Load Balancer and EC2 instances for the application layer in each region
- C. Use a Multi-AZ deployment with MySQL as the data layer.
- D. Use Amazon Route 53 round robin routing to distribute the load evenly to several regions with health check
- E. Host the website on automatically scaled Amazon ECS with AWS Fargate technology containers behind a Network Load Balancer, with an additional Network Load Balancer and Fargate containers for the application layer in each region
- F. Use Amazon Aurora replicas for the data layer.
- G. Use Amazon Route 53 latency-based routing to route to the nearest region with health check
- H. Host the website in Amazon S3 in each region and use Amazon API Gateway with AWS Lambda for the application layer
- I. Use Amazon DynamoDB global tables as the data layer with Amazon DynamoDB Accelerator (DAX) for caching.
- J. Use Amazon Route 53 geolocation-based routing
- K. Host the website on automatically scaled AWS Fargate containers behind a Network Load Balancer with an additional Network Load Balancer and Fargate containers for the application layer in each region
- L. Use Amazon Aurora Multi-Master for Aurora MySQL as the data layer.

**Answer:** C

**Explanation:**

<https://aws.amazon.com/getting-started/projects/build-serverless-web-app-lambda-apigateway-s3-dynamodb-co>

**NEW QUESTION 63**

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

What CI/CD configuration meets all of the requirements?

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

**Answer:** B

**NEW QUESTION 66**

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 70**

A company runs an ordering system on AWS using Amazon SQS and AWS Lambda, with each order received as a JSON message. Recently the company had a marketing event that led to a tenfold increase in orders. With this increase, the following undesired behaviors started in the ordering system:

- Lambda failures while processing orders lead to queue backlogs.
- The same orders have been processed multiple times.

A solutions Architect has been asked to solve the existing issues with the ordering system and add the following resiliency features:

- Retain problematic orders for analysis.
- Send notification if errors go beyond a threshold value. How should the Solutions Architect meet these requirements?

- A. Receive multiple messages with each Lambda invocation, add error handling to message processing code and delete messages after processing, increase the visibility timeout for the messages, create a dead letter queue for messages that could not be processed, create an Amazon CloudWatch alarm on Lambda errors for notification.
- B. Receive single messages with each Lambda invocation, put additional Lambda workers to poll the queue, delete messages after processing, increase the message timer for the messages, use Amazon CloudWatch Logs for messages that could not be processed, create a CloudWatch alarm on Lambda errors for notification.
- C. Receive multiple messages with each Lambda invocation, use long polling when receiving the messages, log the errors from the message processing code using Amazon CloudWatch Logs, create a dead letter queue with AWS Lambda to capture failed invocations, create CloudWatch events on Lambda errors for notification.
- D. Receive multiple messages with each Lambda invocation, add error handling to message processing code and delete messages after processing, increase the visibility timeout for the messages, create a delay queue for messages that could not be processed, create an Amazon CloudWatch metric on Lambda errors for notification.

**Answer:** D

**NEW QUESTION 75**

A Solutions Architect has created an AWS CloudFormation template for a three-tier application that contains an Auto Scaling group of Amazon EC2 instances running a custom AMI.

The Solutions Architect wants to ensure that future updates to the custom AMI can be deployed to a running stack by first updating the template to refer to the new AMI, and then invoking UpdateStack to replace the EC2 instances with instances launched from the new AMI.

How can updates to the AMI be deployed to meet these requirements?

- A. Create a change set for a new version of the template, view the changes to the running EC2 instances to ensure that the AMI is correctly updated, and then execute the change set.
- B. Edit the AWS::AutoScaling::LaunchConfiguration resource in the template, changing its DeletionPolicy to Replace.
- C. Edit the AWS::AutoScaling::AutoScalingGroup resource in the template, inserting an UpdatePolicy attribute.
- D. Create a new stack from the updated template
- E. Once it is successfully deployed, modify the DNS records to point to the new stack and delete the old stack.

**Answer:** C

**Explanation:**

References:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-as-launchconfig.html>

**NEW QUESTION 79**

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 80**

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 84**

A large company experienced a drastic increase in its monthly AWS spend. This is after Developers accidentally launched Amazon EC2 instances in unexpected regions. The company has established practices around least privileges for Developers and controls access to on-premises resources using Active Directory groups. The company now wants to control costs by restricting the level of access that Developers have to the AWS Management Console without impacting their productivity. The company would also like to allow Developers to launch Amazon EC2 in only one region, without limiting access to other services in any region. How can this company achieve these new security requirements while minimizing the administrative burden on the Operations team?

- A. Set up SAML-based authentication tied to an IAM role that has an AdministrativeAccess managed policy attached to it
- B. Attach a customer managed policy that denies access to Amazon EC2 in each region except for the one required.
- C. Create an IAM user for each Developer and add them to the developer IAM group that has the PowerUserAccess managed policy attached to it
- D. Attach a customer managed policy that allows the Developers access to Amazon EC2 only in the required region.
- E. Set up SAML-based authentication tied to an IAM role that has a PowerUserAccess managed policy and a customer managed policy that deny all the Developers access to any AWS services except AWS Service Catalog
- F. Within AWS Service Catalog, create a product containing only the EC2 resources in the approved region.
- G. Set up SAML-based authentication tied to an IAM role that has the PowerUserAccess managed policy attached to it
- H. Attach a customer managed policy that denies access to Amazon EC2 in each region except for the one required.

**Answer:** D

**Explanation:**

The tricks here are: - SAML for AD federation and authentication - PowerUserAccess vs AdministrativeAccess. (PowerUser has less privilege, which is the required one for developers). Admin, has more rights. The description of "PowerUser access" given by AWS is "Provides full access to AWS services and resources, but does not allow management of Users and groups."

**NEW QUESTION 85**

An enterprise runs 103 line-of-business applications on virtual machines in an on-premises data center. Many of the applications are simple PHP, Java, or Ruby web applications, are no longer actively developed, and serve little traffic.

Which approach should be used to migrate these applications to AWS with the LOWEST infrastructure costs ?

- A. Deploy the applications to single-instance AWS Elastic Beanstalk environments without a load balancer.
- B. Use AWS SMS to create AMIs for each virtual machine and run them in Amazon EC2.
- C. Convert each application to a Docker image and deploy to a small Amazon ECS cluster behind an Application Load Balancer.
- D. Use VM Import/Export to create AMIs for each virtual machine and run them in single-instance AWS Elastic Beanstalk environments by configuring a custom image.

**Answer:** A

**Explanation:**

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

**NEW QUESTION 89**

A company has created an account for individual Development teams, resulting in a total of 200 accounts. All accounts have a single virtual private cloud (VPC) in a single region with multiple microservices running in Docker containers that need to communicate with microservices in other accounts. The Security team requirements state that these microservices must not traverse the public internet, and only certain internal services should be allowed to call other individual services. If there is any denied network traffic for a service, the Security team must be notified of any denied requests, including the source IP. How can connectivity be established between services while meeting the security requirements?

- A. Create a VPC peering connection between the VPC
- B. Use security groups on the instances to allow traffic from the security group IDs that are permitted to call the microservice
- C. Apply network ACLs to and allow traffic from the local VPC and peered VPCs only
- D. Within the task definition in Amazon ECS for each of the microservices, specify a log configuration by using the awslogs driver
- E. Within Amazon CloudWatch Logs, create a metric filter and alarm off of the number of HTTP 403 response
- F. Create an alarm when the number of messages exceeds a threshold set by the Security team.
- G. Ensure that no CIDR ranges are overlapping, and attach a virtual private gateway (VGW) to each VPC. Provision an IPsec tunnel between each VGW and enable route propagation on the route table
- H. Configure security groups on each service to allow the CIDR ranges of the VPCs on the other account
- I. Enable VPC Flow Logs, and use an Amazon CloudWatch Logs subscription filter for rejected traffic
- J. Create an IAM role and allow the Security team to call the AssumeRole action for each account.
- K. Deploy a transit VPC by using third-party marketplace VPN appliances running on Amazon EC2, dynamically routed VPN connections between the VPN appliance, and the virtual private gateways (VGWs) attached to each VPC within the region
- L. Adjust network ACLs to allow traffic from the local VPC only
- M. Apply security groups to the microservices to allow traffic from the VPN appliances only
- N. Install the awslogs agent on each VPN appliance, and configure logs to forward to Amazon CloudWatch Logs in the security account for the Security team to access.
- O. Create a Network Load Balancer (NLB) for each microservice
- P. Attach the NLB to a PrivateLink endpoint service and whitelist the accounts that will be consuming this service
- Q. Create an interface endpoint in the consumer VPC and associate a security group that allows only the security group IDs of the services authorized to call the producer service
- R. On the producer services, create security groups for each microservice and allow only the CIDR range of the allowed service
- S. Create VPC Flow Logs on each VPC to capture rejected traffic that will be delivered to an Amazon CloudWatch Logs group
- T. Create a CloudWatch Logs subscription that streams the log data to a security account.

**Answer:** D

**Explanation:**

AWS PrivateLink provides private connectivity between VPCs, AWS services, and on-premises applications, securely on the Amazon network. AWS PrivateLink makes it easy to connect services across different accounts and VPCs to significantly simplify the network architecture. It seems like the next VPC peering.  
<https://aws.amazon.com/privatelink/>

**NEW QUESTION 93**

A company is creating an account strategy so that they can begin using AWS. The Security team will provide each team with the permissions they need to follow the principle of least privileged access. Teams would like to keep their resources isolated from other groups, and the Finance team would like each team's resource usage separated for billing purposes.

Which account creation process meets these requirements and allows for changes?

- A. Create a new AWS Organizations account
- B. Create groups in Active Directory and assign them to roles in AWS to grant federated access
- C. Require each team to tag their resources, and separate bills based on tag
- D. Control access to resources through IAM granting the minimally required privilege.
- E. Create individual accounts for each team
- F. Assign the security as the master account, and enable consolidated billing for all other accounts
- G. Create a cross-account role for security to manage accounts, and send logs to a bucket in the security account.
- H. Create a new AWS account, and use AWS Service Catalog to provide teams with the required resources. Implement a third-party billing to provide the Finance team with the resource use for each team based on tagging
- I. Isolate resources using IAM to avoid account sprawl
- J. Security will control and monitor logs and permissions.
- K. Create a master account for billing using Organizations, and create each team's account from that master account
- L. Create a security account for logs and cross-account access
- M. Apply service control policies on each account, and grant the Security team cross-account access to all accounts
- N. Security will create IAM policies for each account to maintain least privilege access.

**Answer:** B

**NEW QUESTION 95**

A company runs a Windows Server host in a public subnet that is configured to allow a team of administrators to connect over RDP to troubleshoot issues with hosts in a private subnet. The host must be available at all times outside of a scheduled maintenance window, and needs to receive the latest operating system updates within 3 days of release.

What should be done to manage the host with the LEAST amount of administrative effort?

- A. Run the host in a single-instance AWS Elastic Beanstalk environment
- B. Configure the environment with a custom AMI to use a hardened machine image from AWS Marketplace
- C. Apply system updates with AWS Systems Manager Patch Manager.
- D. Run the host on AWS WorkSpace
- E. Use Amazon WorkSpaces Application Manager (WAM) to harden the host
- F. Configure Windows automatic updates to occur every 3 days.
- G. Run the host in an Auto Scaling group with a minimum and maximum instance count of 1. Use a hardened machine image from AWS Marketplace
- H. Apply system updates with AWS Systems Manager Patch Manager.
- I. Run the host in AWS OpsWorks Stack
- J. Use a Chef recipe to harden the AMI during instance launch. Use an AWS Lambda scheduled event to run the Upgrade Operating System stack command to apply system updates.

**Answer: B**

#### NEW QUESTION 96

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

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

- A. Ensure that the Amazon Redshift cluster has been set up to make use of Auto Scaling groups with the nodes in the cluster spread across multiple Availability Zones.
- B. Ensure that the Amazon Redshift cluster creation has been templated using AWS CloudFormation so it can easily be launched in another Availability Zone and data populated from the automated Redshift back-ups stored in Amazon S3.
- C. Use Amazon Kinesis Data Firehose to collect the data ahead of ingestion into Amazon Redshift and create clusters using AWS CloudFormation in another region and stream the data to both clusters.
- D. Create two identical Amazon Redshift clusters in different regions (one as the primary, one as the secondary). Use Amazon S3 cross-region replication from the primary to secondary. Use Amazon S3 cross-region replication from the primary to secondary region, which triggers an AWS Lambda function to populate the cluster in the secondary region.

**Answer: B**

#### Explanation:

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

FROM 37

#### NEW QUESTION 98

A company runs a video processing platform. Files are uploaded by users who connect to a web server, which stores them on an Amazon EFS share. This web server is running on a single Amazon EC2 instance. A different group of instances, running in an Auto Scaling group, scans the EFS share directory structure for new files to process and generates new videos (thumbnails, different resolution, compression, etc.) according to the instructions file, which is uploaded along with the video files. A different application running on a group of instances managed by an Auto Scaling group processes the video files and then deletes them from the EFS share. The results are stored in an S3 bucket. Links to the processed video files are emailed to the customer.

The company has recently discovered that as they add more instances to the Auto Scaling Group, many files are processed twice, so image processing speed is not improved. The maximum size of these video files is 2GB.

What should the Solutions Architect do to improve reliability and reduce the redundant processing of video files?

- A. Modify the web application to upload the video files directly to Amazon S3. Use Amazon CloudWatch Events to trigger an AWS Lambda function every time a file is uploaded, and have this Lambda function put a message into an Amazon SQS queue
- B. Modify the video processing application to read from SQS queue for new files and use the queue depth metric to scale instances in the video processing Auto Scaling group.
- C. Set up a cron job on the web server instance to synchronize the contents of the EFS share into Amazon S3. Trigger an AWS Lambda function every time a file is uploaded to process the video file and store the results in Amazon S3. Using Amazon CloudWatch Events trigger an Amazon SES job to send an email to the customer containing the link to the processed file.
- D. Rewrite the web application to run directly from Amazon S3 and use Amazon API Gateway to upload the video files to an S3 bucket
- E. Use an S3 trigger to run an AWS Lambda function each time a file is uploaded to process and store new video files in a different bucket
- F. Using CloudWatch Events, trigger an SES job to send an email to the customer containing the link to the processed file.
- G. Rewrite the web application to run from Amazon S3 and upload the video files to an S3 bucket
- H. Each time a new file is uploaded, trigger an AWS Lambda function to put a message in an SQS queue containing the link and the instruction
- I. Modify the video processing application to read from the SQS queue and the S3 bucket
- J. Use the queue depth metric to adjust the size of the Auto Scaling group for video processing instances.

**Answer: A**

#### NEW QUESTION 102

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 DM
- 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 AWS
- 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 104

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 106

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 111

An enterprise company is using a multi-account AWS strategy. There are separate accounts for development, staging, and production workloads. To control costs and improve governance, the following requirements have been defined:

- The company must be able to calculate the AWS costs for each project
- The company must be able to calculate the AWS costs for each environment: development, staging, and production
- Commonly deployed IT services must be centrally managed
- Business units can deploy pre-approved IT services only
- Usage of AWS resources in the development account must be limited

Which combination of actions should be taken to meet these requirements? (Select THREE)

- A. Apply environment, cost center, and application name tags to all taggable resources
- B. Configure custom budgets and define thresholds using Cost Explorer
- C. Configure AWS Trusted Advisor to obtain weekly emails with cost-saving estimates
- D. Create a portfolio for each business unit and add products to the portfolios using AWS CloudFormation in AWS Service Catalog
- E. Configure a billing alarm in Amazon CloudWatch.
- F. Configure SCPs in AWS Organizations to allow services available using AWS

**Answer:** CEF

#### NEW QUESTION 112

A company is refactoring an existing web service that provides read and write access to structured data. The service must respond to short but significant spikes in the system load. The service must be fault tolerant across multiple AWS Regions.

Which actions should be taken to meet these requirements?

- A. Store the data in Amazon DocumentDB. Create a single global Amazon CloudFront distribution with a custom origin built on edge-optimized Amazon API Gateway and AWS Lambda. Assign the company's domain as an alternate domain for the distribution.
- B. and configure Amazon Route 53 with an alias to the CloudFront distribution.
- C. Store the data in replicated Amazon S3 buckets in two Regions. Create an Amazon CloudFront distribution in each Region, with custom origins built on Amazon API Gateway and AWS Lambda launched in each Region. Assign the company's domain as an alternate domain for both distributions and configure Amazon Route 53 with a failover routing policy between them.
- D. Store the data in an Amazon DynamoDB global table in two Regions using on-demand capacity mode. In both Regions, run the web service as Amazon ECS Fargate tasks in an Auto Scaling ECS service behind an Application Load Balancer (ALB). In Amazon Route 53, configure an alias record in the company's domain and a Route 53 latency-based routing policy with health checks to distribute traffic between the two ALBs.

Answer: A

#### NEW QUESTION 114

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 failure 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 Redshif
- 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 Forehose delivery strea
- 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-tim
- 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-tim
- 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 strea
- 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 115

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 volumesto 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 119

AnyCompany has acquired numerous companies over the past few years. The CIO for AnyCompany would like to keep the resources for each acquired company separate. The CIO also would like to enforce a chargeback model where each company pays for the AWS services it uses.

The Solutions Architect is tasked with designing an AWS architecture that allows AnyCompany to achieve the following:

- Implementing a detailed chargeback mechanism to ensure that each company pays for the resources it uses.
- AnyCompany can pay for AWS services for all its companies through a single invoice.
- Developers in each acquired company have access to resources in their company only.
- Developers in an acquired company should not be able to affect resources in their company only.
- A single identity store is used to authenticate Developers across all companies.

Which of the following approaches would meet these requirements? (Choose two.)

- A. Create a multi-account strategy with an account per compan
- B. Use consolidated billing to ensure that AnyCompany needs to pay a single bill only.
- C. Create a multi-account strategy with a virtual private cloud (VPC) for each compan
- D. Reduce impact across companies by not creating any VPC peering link
- E. As everything is in a single account, there will be a single invoic
- F. use tagging to create a detailed bill for each company.
- G. Create IAM users for each Developer in the account to which they require acces
- H. Create policies that allow the users access to all resources in that accoun
- I. Attach the policies to the IAM user.
- J. Create a federated identity store against the company's Active Director
- K. Create IAM roles with appropriate permissions and set the trust relationships with AWS and the identity stor
- L. Use AWS STS to grant users access based on the groups they belong to in the identity store.
- M. Create a multi-account strategy with an account per compan
- N. For billing purposes, use a tagging solution that uses a tag to identify the company that creates each resource.

Answer: AD

#### NEW QUESTION 122

A company has more than 100 AWS accounts, with one VPC per account, that need outbound HTTPS connectivity to the internet. The current design contains one

NAT gateway per Availability Zone (AZ) in each VPC. To reduce costs and obtain information about outbound traffic, management has asked for a new architecture for internet access.

Which solution will meet the current needs, and continue to grow as new accounts are provisioned, while reducing costs?

- A. Create a transit VPC across two AZs using a third-party routing appliance
- B. Create a VPN connection to each VPC
- C. Default route internet traffic to the transit VPC.
- D. Create multiple hosted-private AWS Direct Connect VIFs, one per account, each with a Direct Connect gateway
- E. Default route internet traffic back to an on-premises router to route to the internet.
- F. Create a central VPC for outbound internet traffic
- G. Use VPC peering to default route to a set of redundant NAT gateway in the central VPC.
- H. Create a proxy fleet in a central VPC account
- I. Create an AWS PrivateLink endpoint service in the central VPC
- J. Use PrivateLink interface for internet connectivity through the proxy fleet.

**Answer:** D

**Explanation:**

user proxy fleet over PrivateLink. As explained in this AWS website:

<https://aws.amazon.com/blogs/networking-and-content-delivery/how-to-use-aws-privatelink-to-secure-and-scale>

**NEW QUESTION 125**

The company Security team requires that all data uploaded into an Amazon S3 bucket must be encrypted. The encryption keys must be highly available and the company must be able to control access on a per-user basis, with different users having access to different encryption keys.

Which of the following architectures will meet these requirements? (Choose two.)

- A. Use Amazon S3 server-side encryption with Amazon S3-managed key
- B. Allow Amazon S3 to generate an AWS/S3 master key, and use IAM to control access to the data keys that are generated.
- C. Use Amazon S3 server-side encryption with AWS KMS-managed keys, create multiple customer master keys, and use key policies to control access to them.
- D. Use Amazon S3 server-side encryption with customer-managed keys, and use AWS CloudHSM to manage the key
- E. Use CloudHSM client software to control access to the keys that are generated.
- F. Use Amazon S3 server-side encryption with customer-managed keys, and use two AWS CloudHSM instances configured in high-availability mode to manage the key
- G. Use the Cloud HSM client software to control access to the keys that are generated.
- H. Use Amazon S3 server-side encryption with customer-managed keys, and use two AWS CloudHSM instances configured in high-availability mode to manage the key
- I. Use IAM to control access to the keys that are generated in CloudHSM.

**Answer:** BD

**Explanation:**

<http://websecuritypatterns.com/blogs/2018/03/01/encryption-and-key-management-in-aws-kms-vs-cloudhsm-mys/>

**NEW QUESTION 130**

An internal security audit of AWS resources within a company found that a number of Amazon EC2 instances running Microsoft Windows workloads were missing several important operating system-level patches. A Solutions Architect has been asked to fix existing patch deficiencies, and to develop a workflow to ensure that future patching requirements are identified and taken care of quickly. The Solutions Architect has decided to use AWS Systems Manager. It is important that EC2 instance reboots do not occur at the same time on all Windows workloads to meet organizational uptime requirements.

Which workflow will meet these requirements in an automated manner?

- A. Add a Patch Group tag with a value of Windows Servers to all existing EC2 instances
- B. Ensure that all Windows EC2 instances are assigned this tag
- C. Associate the AWS-DefaultPatchBaseline to the Windows servers patch group
- D. Define an AWS Systems Manager maintenance window, conduct patching within it, and associate it with the Windows Servers patch group
- E. Register instances with the maintenance window using associated subnet ID
- F. Assign the AWS-RunPatchBaseline document as a task within each maintenance window.
- G. Add a Patch Group tag with a value of Windows Servers to all existing EC2 instances
- H. Ensure that all Windows EC2 instances are assigned this tag
- I. Associate the AWS-WindowsPatchBaseline document as a task associated with the Windows Servers patch group
- J. Create an Amazon CloudWatch Events rule configured to use a cron expression to schedule the execution of patching using the AWS Systems Manager run command
- K. Create an AWS Systems Manager State Manager document to define commands to be executed during patch execution.
- L. Add a Patch Group tag with a value of either Windows Servers1 or Windows Servers2 to all existing EC2 instances
- M. Ensure that all Windows EC2 instances are assigned this tag
- N. Associate the AWS-DefaultPatchBaseline with both Windows Servers patch groups
- O. Define two non-overlapping AWS Systems Manager maintenance windows, conduct patching within them, and associate each with a different patch group
- P. Register targets with specific maintenance windows using the Patch Group tag
- Q. Assign the AWS-RunPatchBaseline document as a task within each maintenance window.
- R. Add a Patch Group tag with a value of either Windows Servers1 or Windows Servers2 to all existing EC2 instances
- S. Ensure that all Windows EC2 instances are assigned this tag
- T. Associate the AWS-WindowsPatchBaseline with both Windows Servers patch groups
- . Define two non-overlapping AWS Systems Manager maintenance windows, conduct patching within them, and associate each with a different patch group
- . Assign the AWS-RunWindowsPatchBaseline document as a task within each maintenance window
- . Create an AWS Systems Manager State Manager document to define commands to be executed during patch execution.

**Answer:** C

**NEW QUESTION 135**

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 138

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 143

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 146

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 ProvisionedThrough put Exceeded 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 loa
- 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 151

A company is planning to migrate an application from on-premises to AWS. The application currently uses an Oracle database and the company can tolerate a brief downtime of 1 hour when performing the switch to the new infrastructure. As part of the migration, the database engine will be changed to MySQL. A Solutions Architect needs to determine which AWS services can be used to perform the migration while minimizing the amount of work and time required.

Which of the following will meet the requirements?

- A. Use AWS SCT to generate the schema scripts and apply them on the target prior to migratio
- B. Use AWS DMS to analyse the current schema and provide a recommendation for the optimal database engin
- C. Then, use AWS DMS to migrate to the recommended enginee
- D. Use AWS SCT to identify what embedded SQL code in the application can be converted and what has to be done manually.
- E. Use AWS SCT to generate the schema scripts and apply them on the target prior to migratio
- F. Use AWS DMS to begin moving data from the on-premises database to AW
- G. After the initial copy, continue to use AWS DMS to keep the databases in sync until cutting over to the new databas
- H. Use AWS SCT to identify what embedded SQL code in the application can be converted and what has to be done manually.
- I. Use AWS DMS to help identify the best target deployment between installing the database engine on Amazon EC2 directly or moving to Amazon RD
- J. Then, use AWS DMS to migrate to the platfor
- K. Use AWS Application Discovery Service to identify what embedded SQL code in the application can be converted and what has to be done manually.
- L. Use AWS DMS to begin moving data from the on-premises database to AW
- M. After the initial copy, continue to use AWS DMS to keep the databases in sync until cutting over to the new databas
- N. Use AWS Application Discovery Service to identify what embedded SQL code in the application can be converted and what has to be done manually.

**Answer: B**

#### NEW QUESTION 154

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

Which solution will meet the CISO's requirements?

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

**Answer: A**

#### Explanation:

[https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial\\_cross-account-with-roles.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html)

#### NEW QUESTION 155

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 157

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

- \* The application code will need a 500 HB 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 unpatche
- B. Use AWS CodeDeploy to push the application code to the instance
- C. Store the static data in Amazon EFS.
- D. Use AWS System Manager to create a new AMI with upload OS patche
- E. Update the Auto Scaling group to use the patches AMI and replace existing unpatches and the application code as a batch job every nigh
- 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 coun
- H. Configure an Amazon EC2 data script to download the data from Amazon S3 install OS patches with AWS system Manager when they are release
- I. Use Codedeploy to push the application code to the instances.

- J. Use an Amazon provided AMI for the OS Configure an Auto Scaling group Configure an Amazon EC2 user data script to download the data from Amazon S3.  
 Replace existing instances after each Amazon-provided AMI releases
- K. Use AWS CodeDeploy to push the application code to the instances.

**Answer: C**

#### NEW QUESTION 158

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 retrieve data to be processed.
  - Grid instances must communicate with Amazon DynamoDB to track intermediate data,
  - The job scheduler need 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 163

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 167

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 VP
- 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 if 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 169

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 172

A company currently runs a secure application on Amazon EC2 that takes files from on-premises locations through AWS Direct Connect, processes them, and uploads them to a single Amazon S3 bucket. The application uses HTTPS for encryption in transit to Amazon S3, and S3 server-side encryption to encrypt at rest. Which of the following changes should the Solutions Architect recommend to make this solution more secure without impeding application's performance?

- A. Add a NAT gateway
- B. Update the security groups on the EC2 instance to allow access to and from the S3 IP range only
- C. Configure an S3 bucket policy that allows communication from the NAT gateway's Elastic IP address only.
- D. Add a VPC endpoint
- E. Configure endpoint policies on the VPC endpoint to allow access to the required Amazon S3 buckets only
- F. Implement an S3 bucket policy that allows communication from the VPC's source IP range only.
- G. Add a NAT gateway
- H. Update the security groups on the EC2 instance to allow access to and from the S3 IP range only
- I. Configure an S3 bucket policy that allows communication from the source public IP address of the on-premises network only.
- J. Add a VPC endpoint
- K. Configure endpoint policies on the VPC endpoint to allow access to the required S3 buckets only
- L. Implement an S3 bucket policy that allows communication from the VPC endpoint only.

**Answer:** D

#### Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/dev/example-bucket-policies-vpc-endpoint.html>

#### NEW QUESTION 175

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

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

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

**Answer:** BDE

#### NEW QUESTION 176

A Solutions Architect needs to design a highly available application that will allow authenticated users to stay connected to the application even when there are underlying failures

Which solution will meet these requirements?

- A. Deploy the application on Amazon EC2 instances Use Amazon Route 53 to forward requests to the EC2 Instances Use Amazon DynamoDB to save the authenticated connection details
- B. Deploy the application on Amazon EC2 instances in an Auto Scaling group Use an internet-facing Application Load Balancer to handle requests Use Amazon DynamoDB to save the authenticated connection details
- C. Deploy the application on Amazon EC2 instances in an Auto Scaling group Use an internet-facing Application Load Balancer on the front end Use EC2 instances to save the authenticated connection details
- D. Deploy the application on Amazon EC2 instances in an Auto Scaling group Use an internet-facing Application Load Balancer on the front end Use EC2 instances hosting a MySQL database to save the authenticated connection details

**Answer:** B

#### NEW QUESTION 178

A large company is migrating its entire IT portfolio to AWS. Each business unit in the company has a standalone AWS account that supports both development and test environments. New accounts to support production workloads will be needed soon.

The Finance department requires a centralized method for payment but must maintain visibility into each group's spending to allocate costs.

The Security team requires a centralized mechanism to control IAM usage in all the company's accounts. What combination of the following options meet the company's needs with LEAST effort? (Choose two.)

- A. Use a collection of parameterized AWS CloudFormation templates defining common IAM permissions that are launched into each account
- B. Require all new and existing accounts to launch the appropriate stacks to enforce the least privilege model.
- C. Use AWS Organizations to create a new organization from a chosen payer account and define an organizational unit hierarchy
- D. Invite the existing accounts to join the organization and create new accounts using Organizations.
- E. Require each business unit to use its own AWS account
- F. Tag each AWS account appropriately and enable Cost Explorer to administer chargebacks.
- G. Enable all features of AWS Organizations and establish appropriate service control policies that filter IAM permissions for sub-accounts.
- H. Consolidate all of the company's AWS accounts into a single AWS account
- I. Use tags for billing purposes and IAM's Access Advice feature to enforce the least privilege model.

**Answer:** BD

#### NEW QUESTION 181

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 184**

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

### Money Back Guarantee

#### **SAP-C01 Practice Exam Features:**

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