

Amazon-Web-Services

Exam Questions SAP-C01

AWS Certified Solutions Architect- Professional



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

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 4

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

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

- N. Generate presigned URLs when returning references to content stored in Amazon S3.
- O. Set up Amazon API Gateway and create the required API resources and method
- P. Use an Amazon API Gateway custom authorizer to control access to the AP
- Q. Configure the methods to use AWS Lambda custom integrations, and process each resource with a unique Lambda functio
- R. Store request meta information in an Amazon ElastiCache Multi-AZ cluster and static content in a secured S3 bucke
- S. Generate presigned URLs when returning references to content stored in Amazon S3.

Answer: C

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.

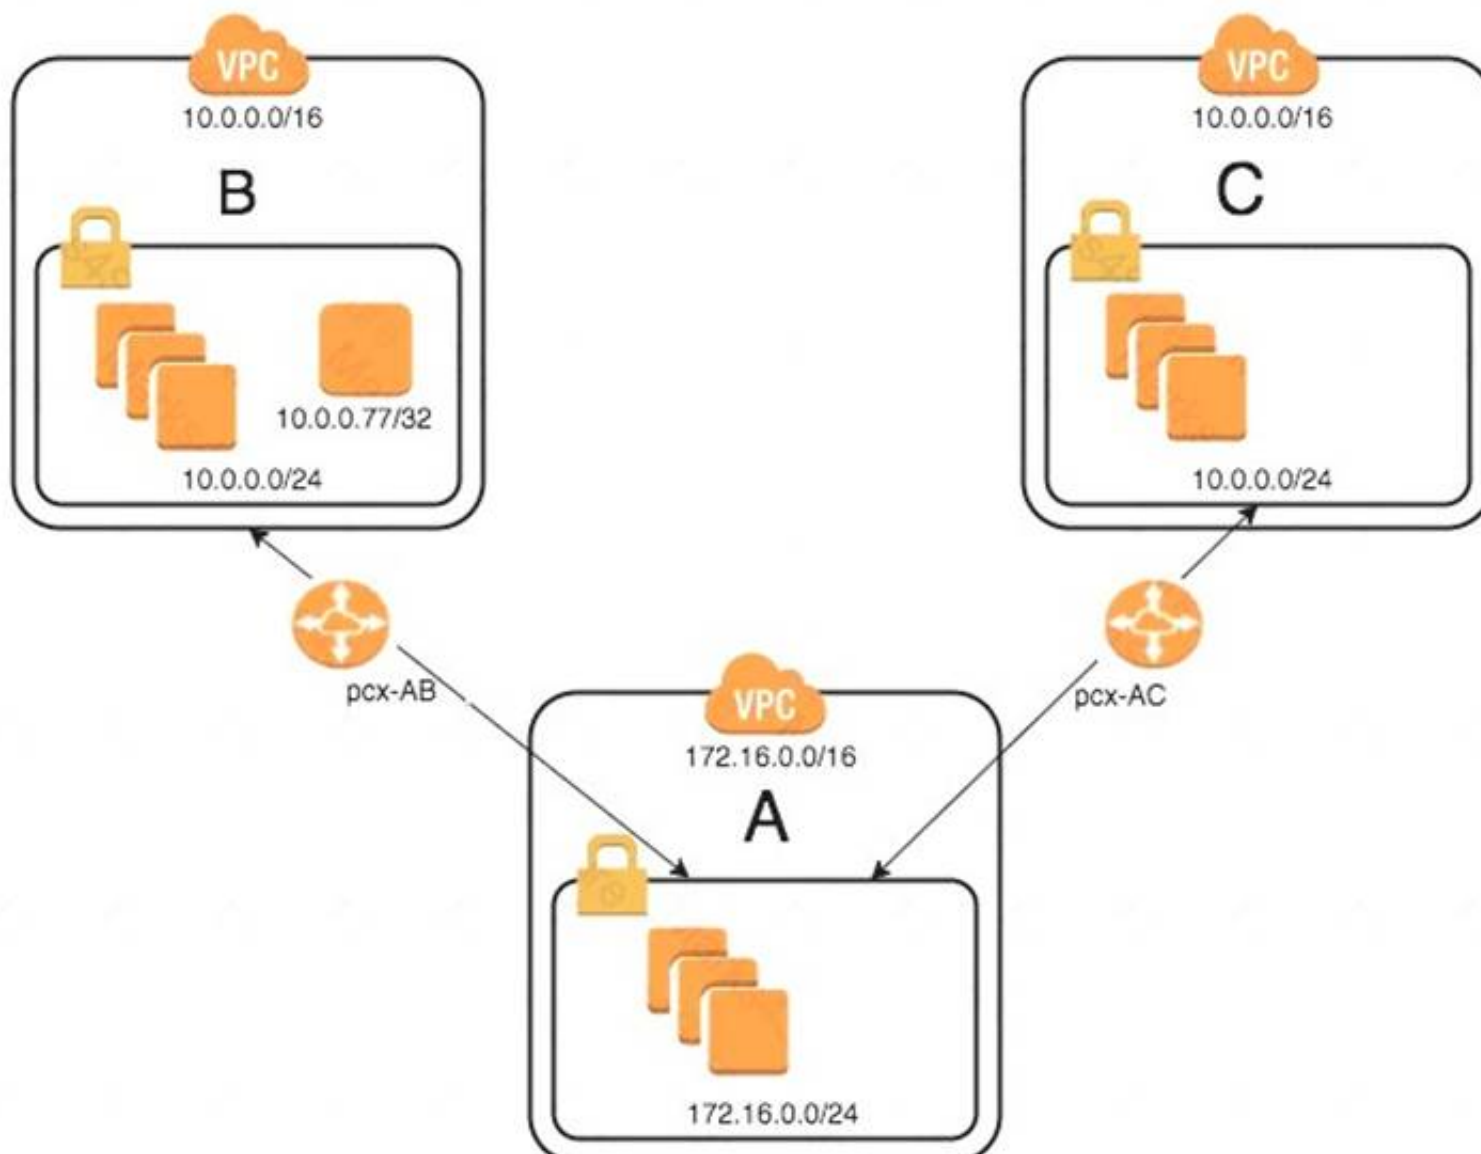
How can the design requirements be met?

- A. Use Amazon CloudFront with Amazon ECS for hosting the websit
- 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 websit
- 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

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 7

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

NEW QUESTION 8

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

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

Answer: D

NEW QUESTION 9

A company uses an Amazon EMR cluster to process data once a day. The raw data comes from Amazon S3, and the resulting processed data is also stored in Amazon S3. The processing must complete within 4 hours; currently, it only takes 3 hours. However, the processing time is taking 5 to 10 minutes longer each week due to an increasing volume of raw data. The team is also concerned about rising costs as the compute capacity increases. The EMR cluster is currently running on three m3.xlarge instances (one master and two core nodes). Which of the following solutions will reduce costs related to the increasing compute needs?

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

node

G. Purchase a standard all-upfront Reserved Instance for the master node.

Answer: B

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

The following is an example of the additional data:

```
list celebrities [name of the personality] wearing [color] looking [happy, sad] near [location example Eiffel Tower in Paris]
```

As part of the cloud migration program, the company uploaded existing image data to Amazon S3 and told users to upload images directly to Amazon S3.

What should the Solutions Architect do to support these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 15

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 functio
- 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 cod
- D. Rollback if Amazon CloudWatch alarms are triggered.
- E. Refactor the AWS CLI scripts into a single script that deploys the new Lambda versio
- F. When deployment is completed, the script tests execut
- 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 versio
- 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-serverle>

NEW QUESTION 17

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 19

A company with 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 20

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 25

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 configuration.
- D. Host the application and presentation layers in an Auto Scaling configuration on Amazon EC2 instances behind an Application Load Balance.

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

Answer: B

NEW QUESTION 30

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

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

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

Answer: C

NEW QUESTION 31

A development team has created a series of AWS CloudFormation templates to help deploy services. They created a template for a network/virtual private (VPC) stack, a database stack, a bastion host stack, and a web application-specific stack. Each service requires the deployment of at least:

Each template has multiple input parameters that make it difficult to deploy the services individually from the AWS CloudFormation console. The input parameters from one stack are typically outputs from other stacks. For example, the VPC ID, subnet IDs, and security groups from the network stack may need to be used in the application stack or database stack.

Which actions will help reduce the operational burden and the number of parameters passed into a service deployment? (Choose two.)

- A. Create a new AWS CloudFormation template for each servic
- B. After the existing templates to use cross-stack references to eliminate passing many parameters to each templat
- C. Call each required stack for the application as a nested stack from the new stac
- D. Call the newly created service stack from the AWS CloudFormation console to deploy the specific service with a subset of the parameters previously required.
- E. Create a new portfolio in AWS Service Catalog for each servic
- F. Create a product for each existing AWS CloudFormation template required to build the servic
- G. Add the products to the portfolio that represents that service in AWS Service Catalo
- H. To deploy the service, select the specific service portfolio and launch the portfolio with the necessary parameters to deploy all templates.
- I. Set up an AWS CodePipeline workflow for each servic
- J. For each existing template, choose AWS CloudFormation as a deployment actio
- K. Add the AWS CloudFormation template to the deployment actio
- L. Ensure that the deployment actions are processed to make sure that dependences are obeye
- M. Use configuration files and scripts to share parameters between the stack
- N. To launch the service, execute the specific template by choosing the name of the service and releasing a change.
- O. Use AWS Step Functions to define a new servic
- P. Create a new AWS CloudFormation template for each servic
- Q. After the existing templates to use cross-stack references to eliminate passing many parameters to each templat
- R. Call each required stack for the application as a nested stack from the new service templat
- S. Configure AWS Step Functions to call the service template directl
- T. In the AWS Step Functions console, execute the step.
- . Create a new portfolio for the Services in AWS Service Catalo
- . Create a new AWS CloudFormation template for each servic
- . After the existing templates to use cross-stack references to eliminate passing many parameters to each templat
- . Call each required stack for the application as a nested stack from the new stac
- . Create a product for each applicatio
- . Add the service template to the produc
- . Add each new product to the portfoli
- . Deploy the product from the portfolio to deploy the service with the necessary parameters only to start the deployment.

Answer: AE

NEW QUESTION 36

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

Which steps should a Solutions Architect take to reduce costs?

- A. Enable AWS Business Support and review AWS Trusted Advisor's cost check
- B. Create Amazon EC2 Auto Scaling groups for applications that experience fluctuating deman
- C. Save AWS Simple Monthly Calculator reports in Amazon S3 for trend analysi
- D. Create a master account under Organizations and have teams join for consolidating billing.
- E. Enable Cost Explorer and AWS Business Support Reserve Amazon EC2 and Amazon RDS DB instance
- F. Use Amazon CloudWatch and AWS Trusted Advisor for monitoring and to receive cost-savings suggestion

- G. Create a master account under Organizations and have teams join for consolidated billing.
- H. Create an AWS Lambda function that changes the instance size based on Amazon CloudWatch alarms. Reserve instances based on AWS Simple Monthly Calculator suggestion
- I. Have an 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: B

Explanation:

Import/Export supports importing and exporting data into and out of Amazon S3 buckets. For significant data sets, AWS Import/Export is often faster than Internet transfer and more cost effective than upgrading your connectivity.

NEW QUESTION 41

A company runs a public-facing application that uses a Java-based web service via a RESTful API. It is hosted on Apache Tomcat on a single server in a data center that runs consistently at 30% CPU utilization. Use of the API is expected to increase by 10 times with a new product launch. The business wants to migrate the application to AWS with no disruption and needs it to scale to meet demand. The company has already decided to use Amazon Route 53 and CNAME records to redirect traffic. How can these requirements be met with the LEAST amount of effort?

- A. Use AWS Elastic Beanstalk to deploy the Java web service and enable Auto Scaling. Then switch the application to use the new web service.
- B. Lift and shift the Apache server to the cloud using AWS SMS. Then switch the application to direct web service traffic to the new instance.
- C. Create a Docker image and migrate the image to Amazon ECS. Then change the application code to direct web service queries to the ECS container.
- D. Modify the application to call the web service via Amazon API Gateway. Then create a new AWS Lambda Java function to run the Java web service code. After testing, change API Gateway to use the Lambda function.

Answer: A

NEW QUESTION 43

A company is running a .NET three-tier web application on AWS. The team currently uses EBS storage optimized instances to store and 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 servers.
- B. Create an Elastic Load Balancer with Auto Scaling general purpose instances.
- C. Enable Amazon CloudFront to the Elastic Load Balancer.
- 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 Balancer.
- 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 requirements.
- G. Use Cost Explorer and AWS Trusted Advisor checks to continue monitoring the environment for future savings.
- H. Move the entire website to Amazon S3 using the S3 website hosting feature.
- I. Remove all the web servers and have Amazon S3 communicate directly with the application servers in Amazon VPC.
- J. Use AWS Elastic Beanstalk to deploy the .NET application.
- K. Move all images and video files to Amazon EFS.
- L. Create an Amazon CloudFront distribution that points to the EFS share.
- M. Reserve the m4.xlarge instances needed to meet base performance requirements.

Answer: B

NEW QUESTION 46

A company has an Amazon VPC that is divided into a public subnet and a private subnet. A web application runs in Amazon VPC, and each subnet has its own NACL. The public subnet has a CIDR of 10.0.0.0/24. An Application Load Balancer is deployed to the public subnet. The private subnet has a CIDR of 10.0.1.0/24. Amazon EC2 instances that run a web server on port 80 are launched into the private subnet. Only network traffic that is required for the Application Load Balancer to access the web application can be allowed to travel between the public and private subnets. What collection of rules should be written to ensure that the private subnet's NACL meets the requirement? (Select TWO.)

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

Answer: BE

NEW QUESTION 47

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 accounts.

- 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 52

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

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

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

Answer: D

NEW QUESTION 53

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 57

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 mod
- 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 DynamoD
- E. Use an AWS KMS key to encrypt DynamoDB objects at rest.
- F. Deploy instances with Amazon EBS volumes attached to store this dat
- 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 62

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

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

Answer: B

Explanation:

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

NEW QUESTION 64

A retail company has a custom NET web application running on AWS that uses Microsoft SQL Server for the database The application servers maintain a user's session locally.

Which combination of architecture changes are needed ensure all tiers of the solution are highly available? (Select THREE.)

- A. Refactor the application to store the user's session in Amazon ElastiCache Use Application Load Balancers to distribute the load between application instances
- B. Set up the database to generate hourly snapshots using Amazon EBS Configure an Amazon CloudWatch Events rule to launch a new database instance if the primary one fails
- C. Migrate the database to Amazon RDS tor SQL Server Configure the RDS instance to use a Multi-AZ deployment
- D. Move the NET content to an Amazon S3 bucket Configure the bucket for static website hosting
- E. Put the application instances in an Auto Scaling group Configure the Auto Scaling group to create new instances if an instance becomes unhealthy
- F. Deploy Amazon CloudFront in front of the application tier Configure CloudFront to serve content from healthy application instances only

Answer: BDE

NEW QUESTION 68

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

The company wants to ensure that the application remains available even in the event of a region-level failure of all of the application's components. It is

acceptable for the application to be in read-only mode for up to 1 hour. The company plans to configure two Amazon Route 53 record sets, one for each of the regions.

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

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

Answer: CE

Explanation:

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

NEW QUESTION 74

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

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 77

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

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

Answer: BD

Explanation:

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

NEW QUESTION 82

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 denies 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

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

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

Which design meets these requirements?

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

Answer: D

NEW QUESTION 87

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 processes. 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 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 91

A company has a High Performance Computing (HPC) cluster in its on-premises data center which runs thousands of jobs in parallel for one week every month, processing petabytes of images. The images are stored on a network file server, which is replicated to a disaster recovery site. The on-premises data center has reached capacity and has started to spread the jobs out over the course of month in order to better utilize the cluster, causing a delay in the job completion.

The company has asked its Solutions Architect to design a cost-effective solution on AWS to scale beyond the current capacity of 5,000 cores and 10 petabytes of data. The solution must require the least amount of management overhead and maintain the current level of durability.

Which solution will meet the company's requirements?

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

Answer: B

NEW QUESTION 93

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

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 96

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

What is the LEAST complex method of migrating the database securely and reliably?

- A. Order an AWS Snowball device and copy the database using the AWS Data Migration Agent
- B. When the database is available in Amazon S3, use AWS DMS to load it to Amazon RDS, and configure a job to synchronize changes before the cutover.
- C. Create an AWS DMS job to continuously replicate the data from on premises to 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 101

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 103

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

What is the MOST cost-effective solution?

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

Answer: C

NEW QUESTION 106

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 107

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 112

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

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

- A. Use AWS Application Discovery Service and deploy the data collection agent to each virtual machine in the data center.
- B. Configure the Amazon CloudWatch agent on all servers within the local environment and publish metrics to Amazon CloudWatch Logs.
- C. Use AWS Application Discovery Service and enable agentless discovery in the existing virtualization environment.
- D. Enable AWS Application Discovery Service in the AWS Management Console and configure the corporate firewall to allow scans over a VPN.

Answer: A

NEW QUESTION 117

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

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

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

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

- B. Run backups nightly and store the virtual tapes on Amazon S3 standard storage in US-EAST-1. Use cross-region replication to create a second copy in US-WEST-2. Use Amazon S3 lifecycle policies to perform automatic migration to Amazon Glacier and deletion of expired backups after 1 year?
- C. Configure the local source code repository to synchronize files to an AWS Storage Gateway file Amazon gateway to store backup copies in an Amazon S3 Standard bucket
- D. Enable versioning on the Amazon S3 bucket
- E. Create Amazon S3 lifecycle policies to automatically migrate old versions of objects to Amazon S3 Standard Infrequent Access, then Amazon Glacier, then delete backups after 1 year.
- F. Replace the local source code repository storage with a Storage Gateway stored volume
- G. Change the default snapshot frequency to 1 hour
- H. Use Amazon S3 lifecycle policies to archive snapshots to Amazon Glacier and remove old snapshots after 1 year
- I. Use cross-region replication to create a copy of the snapshots in US-WEST-2.
- J. Replace the local source code repository storage with a Storage Gateway cached volume
- K. Create a snapshot schedule to take hourly snapshots
- L. Use an Amazon CloudWatch Events schedule expression rule to run on hourly AWS Lambda task to copy snapshots from US-EAST-1 to US-WEST-2.

Answer: B

Explanation:

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

NEW QUESTION 121

A company manages more than 200 separate internet-facing web applications. All of the applications are deployed to AWS in a single AWS Region. The fully qualified domain names (FQDNs) of all of the applications are made available through HTTPS using Application Load Balancers (ALBs). The ALBs are configured to use public SSL/TLS certificates.

A Solutions Architect needs to migrate the web applications to a multi-region architecture. All HTTPS services should continue to work without interruption. Which approach meets these requirements?

- A. Request a certificate for each FQDN using AWS KMS
- B. Associate the certificates with the ALBs in the primary AWS Region
- C. Enable cross-region availability in AWS KMS for the certificates and associate the certificates with the ALBs in the secondary AWS Region.
- D. Generate the key pairs and certificate requests for each FQDN using AWS KMS
- E. Associate the certificates with the ALBs in both the primary and secondary AWS Regions.
- F. Request a certificate for each FQDN using AWS Certificate Manager
- G. Associate the certificates with the ALBs in both the primary and secondary AWS Regions.
- H. Request certificates for each FQDN in both the primary and secondary AWS Regions using AWS Certificate Manager
- I. Associate the certificates with the corresponding ALBs in each AWS Region.

Answer: D

Explanation:

<https://docs.aws.amazon.com/acm/latest/userguide/acm-regions.html>

Certificates in ACM are regional resources. To use a certificate with Elastic Load Balancing for the same fully qualified domain name (FQDN) or set of FQDNs in more than one AWS region, you must request or import a certificate for each region. For certificates provided by ACM, this means you must revalidate each domain name in the certificate for each region. You cannot copy a certificate between regions.

NEW QUESTION 123

A company's data center is connected to the AWS Cloud over a minimally used 10-Gbps AWS Direct Connect connection with a private virtual interface to its virtual private cloud (VPC). The company internet connection is 200 Mbps and the company has a 150-TB dataset that is created each Friday. The data must be transferred and available in Amazon S3 on Monday morning.

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

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

Answer: D

NEW QUESTION 128

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

Which option will meet these requirements with MINIMAL effort?

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

Answer: B

Explanation:

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

NEW QUESTION 133

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 138

A company has an internal AWS Elastic Beanstalk worker environment inside a VPC that must access an external payment gateway API available on an HTTPS endpoint the public internet. Because of security policies, the payment gateway's Application team can grant access to only one public IP address.

Which architecture will set up an Elastic Beanstalk environment to access the company's application without making multiple changes on the company's end?

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

Answer: C

NEW QUESTION 141

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 144

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

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

The Chief Technology Officer wants to improve the user experience.

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

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

Answer: AE

NEW QUESTION 146

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

Key requirements are:

- Grid instances must communicate with Amazon S3 to retrieve data to be processed.
- Grid instances must communicate with Amazon DynamoDB to track intermediate data.
- The job scheduler 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 149

To abide by industry regulations, a Solutions Architect must design a solution that will store a company's critical data in multiple public AWS Regions, including in the United States, where the company's headquarters is located. The Solutions Architect is required to provide access to the data stored in AWS to the company's global WAN network. The Security team mandates that no traffic accessing this data should traverse the public internet. How should the Solutions Architect design a highly available solution that meets the requirements and is cost-effective?

- A. Establish AWS Direct Connect connections from the company headquarters to all AWS Regions in use. Use the company WAN to send traffic over to the headquarters and then to the respective DX connection to access the data.
- B. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection.
- C. Use inter-region VPC peering to access the data in other AWS Regions.
- D. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection.
- E. Use an AWS transit VPC solution to access data in other AWS Regions.
- F. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection.
- G. Use Direct Connect Gateway to access data in other AWS Regions.

Answer: D

Explanation:

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

NEW QUESTION 150

An online retailer needs to regularly process large product catalogs, which are handled in batches. These are sent out to be processed by people using the Amazon Mechanical Turk service, but the retailer has asked its Solutions Architect to design a workflow orchestration system that allows it to handle multiple concurrent Mechanical Turk operations, deal with the result assessment process, and reprocess failures.

Which of the following options gives the retailer the ability to interrogate the state of every workflow with the LEAST amount of implementation effort?

- A. Trigger Amazon CloudWatch alarms based upon message visibility in multiple Amazon SQS queues (one queue per workflow stage) and send messages via Amazon SNS to trigger AWS Lambda functions to process the next step.
- B. Use Amazon ES and Kibana to visualize Lambda processing logs to see the workflow states.
- C. Hold workflow information in an Amazon RDS instance with AWS Lambda functions polling RDS for status change.
- D. Worker Lambda functions then process the next workflow step.
- E. Amazon QuickSight will visualize workflow states directly out of Amazon RDS.
- F. Build the workflow in AWS Step Functions, using it to orchestrate multiple concurrent workflow.
- G. The status of each workflow can be visualized in the AWS Management Console, and historical data can be written to Amazon S3 and visualized using Amazon QuickSight.
- H. Use Amazon SWF to create a workflow that handles a single batch of catalog records with multiple worker tasks to extract the data, transform it, and send it through Mechanical Turk.
- I. Use Amazon ES and Kibana to visualize AWS Lambda processing logs to see the workflow states.

Answer: C

Explanation:

AWS Step Functions is a fully managed service that makes it easy to coordinate the components of distributed applications and microservices using visual workflows. Instead of writing a Decider program, you define state machines in JSON. AWS customers should consider using Step Functions for new applications. If Step Functions does not fit your needs, then you should consider Amazon Simple Workflow (SWF). Amazon SWF provides you complete control over your orchestration logic, but increases the complexity of developing applications. You may write decider programs in the programming language of your choice, or you may use the Flow framework to use programming constructs that structure asynchronous interactions for you. AWS will continue to provide the Amazon SWF service, Flow framework, and support all Amazon SWF customers. <https://aws.amazon.com/swf/faqs/>

NEW QUESTION 151

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

- Consolidate all accounts into one organization.
- Allow full access to the Amazon EC2 service from the master account and the secondary accounts.
- Minimize the effort required to add additional secondary accounts.

Which combination of steps should be included in the solution? (Choose two.)

- A. Create an organization from the master account
- B. Send invitations to the secondary accounts from the master account
- C. Accept the invitations and create an OU.
- D. Create an organization from the master account
- E. Send a join request to the master account from each secondary account
- F. Accept the requests and create an OU.
- G. Create a VPC peering connection between the master account and the secondary account

- H. Accept the request for the VPC peering connection.
- I. Create a service control policy (SCP) that enables full EC2 access, and attach the policy to the OU.
- J. Create a full EC2 access policy and map the policy to a role in each account
- K. Trust every other account to assume the role.

Answer: AD

Explanation:

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

NEW QUESTION 153

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 154

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 156

A company has a web application that securely uploads pictures and videos to an Amazon S3 bucket The company requires that only authenticated users are allowed to post content The application generates a presigned URL that is used to upload objects through a browser interface Most users are reporting slow upload times for objects larger than 100 MB.

What can a Solutions Architect do to improve the performance of these uploads while ensuring only authenticated users are allowed to post content?

- A. Set up an Amazon API Gateway with an edge-optimized API endpoint that has a resource as an S3 service proxy Configure the PUT method for this resource to expose the S3 Putobject operation Secure the API Gateway using a COGNITO_USER_POOLS authorize
- B. Have the browser interface use API Gateway instead of the presigned URL to upload objects
- C. Set up an Amazon API Gateway with a regional API endpoint that has a resource as an S3 service proxy Configure the PUT method for this resource to expose the S3 Putobject operation Secure the API Gateway using an AWS Lambda authorizer Have the browser interface use API Gateway instead of the presigned URL to upload objects
- D. Enable an S3 Transfer Acceleration endpoint on the S3 bucket Use the endpoint when generating the presigned URL Have the browser interface upload the objects to the URL using the S3 multipart upload API.
- E. Configure an Amazon CloudFront distribution for the destination S3 bucket Enable PUT and POST methods for the CloudFront cache behavior Update the CloudFront origin to use an origin access identity (OAI). Give the OAI user s3:PutObject permissions in the bucket policy Have the browser interface upload objects using the CloudFront distribution.

Answer: A

NEW QUESTION 158

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

Which application migration strategy meets this requirement?

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

Answer: B

Explanation:

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

NEW QUESTION 163

A company is building an AWS landing zone and has asked a Solutions Architect to design a multi-account access strategy that will allow hundreds of users to use corporate credentials to access the AWS Console. The company is running a Microsoft Active Directory and users will use an AWS Direct Connect connection to connect to AWS. The company also wants to be able to federate to third-party services and providers, including custom applications. Which solution meets the requirements by using the LEAST amount of management overhead?

- A. Connect the Active Directory to AWS by using single sign-on and an Active Directory Federation Services (AD FS) with SAML 2.0, and then configure the identity Provider (IdP) system to use form-based authentication
- B. Build the AD FS portal page with corporate branding, and integrate third-party applications that support SAML 2.0 as required.
- C. Create a two-way Forest trust relationship between the on-premises Active Directory and the AWS Directory Service
- D. Set up AWS Single Sign-On with AWS Organization
- E. Use single sign-on integrations for connections with third-party applications.
- F. Configure single sign-on by connecting the on-premises Active Directory using the AWS Directory Service AD Connector
- G. Enable federation to the AWS services and accounts by using the IAM applications and services linking function
- H. Leverage third-party single sign-on as needed.
- I. Connect the company's Active Directory to AWS by using AD FS and SAML 2.0. Configure the AD FS claim rule to leverage Regex and a common Active Directory naming convention for the security group to allow federation of all AWS accounts
- J. Leverage third-party single sign-on as needed, and add it to the AD FS server.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/security/aws-federated-authentication-with-active-directory-federation-services-a>

NEW QUESTION 167

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

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

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

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

Answer: B

NEW QUESTION 172

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

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

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

Answer: AB

NEW QUESTION 176

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

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

What changes should make the bid processing more reliable?

- A. Refactor the web application to use the Amazon Kinesis Producer Library (KPL) when posting bids to Kinesis Data Stream
- B. Refactor the bid processor to flag each record in Kinesis Data Streams as being unread, processing, and processed
- C. At the start of each bid processing run, scan Kinesis Data Streams for unprocessed records.
- D. Refactor the web application to post each incoming bid to an Amazon SNS topic in place of Kinesis Data Stream
- E. Configure the SNS topic to trigger an AWS Lambda function that processes each bid as soon as a user submits it.
- F. Refactor the web application to post each incoming bid to an Amazon SQS FIFO queue in place of Kinesis Data Stream

- G. Refactor the bid processor to continuously the SQS queue
- H. Place the bid processing EC2 instance in an Auto Scaling group with a minimum and a maximum size of 1.
- I. Switch the EC2 instance type from t2.large to a larger general compute instance type
- J. Put the bid processor EC2 instances in an Auto Scaling group that scales out the number of EC2 instances running the bid processor, based on the IncomingRecords metric in Kinesis Data Streams.

Answer: C

Explanation:

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

NEW QUESTION 179

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

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

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

Answer: AEF

NEW QUESTION 184

A company is running multiple applications on Amazon EC2. Each application is deployed and managed by multiple business units. All applications are deployed on a single AWS account but on different virtual private clouds (VPCs). The company uses a separate VPC in the same account for test and development purposes.

Production applications suffered multiple outages when users accidentally terminated and modified resources that belonged to another business unit. A Solutions Architect has been asked to improve the availability of the company applications while allowing the Developers access to the resources they need.

Which option meets the requirements with the LEAST disruption?

- A. Create an AWS account for each business unit
- B. Move each business unit's instances to its own account and set up a federation to allow users to access their business unit's account.
- C. Set up a federation to allow users to use their corporate credentials, and lock the users down to their own VPC
- D. Use a network ACL to block each VPC from accessing other VPCs.
- E. Implement a tagging policy based on business unit
- F. Create an IAM policy so that each user can terminate instances belonging to their own business units only.
- G. Set up role-based access for each user and provide limited permissions based on individual roles and the services for which each user is responsible.

Answer: C

Explanation:

Principal – Control what the person making the request (the principal) is allowed to do based on the tags that are attached to that person's IAM user or role. To do this, use the aws:PrincipalTag/key-name condition key to specify what tags must be attached to the IAM user or role before the request is allowed.

https://docs.aws.amazon.com/IAM/latest/UserGuide/access_iam-tags.html

NEW QUESTION 189

A company has released a new version of a website to target an audience in Asia and South America. The website's media assets are hosted on Amazon S3 and have an Amazon CloudFront distribution to improve end-user performance. However, users are having a poor login experience the authentication service is only available in the us-east-1 AWS Region.

How can the Solutions Architect improve the login experience and maintain high security and performance with minimal management overhead?

- A. Replicate the setup in each new geography and use Amazon Route 53 geo-based routing to route traffic to the AWS Region closest to the users.
- B. Use an Amazon Route 53 weighted routing policy to route traffic to the CloudFront distributio
- C. Use CloudFront cached HTTP methods to improve the user login experience.
- D. Use Amazon Lambda@Edge attached to the CloudFront viewer request trigger to authenticate and authorize users by maintaining a secure cookie token with a session expiry to improve the user experience in multiple geographies.
- E. Replicate the setup in each geography and use Network Load Balancers to route traffic to the authentication service running in the closest region to users.

Answer: C

Explanation:

There are several benefits to using Lambda@Edge for authorization operations. First, performance is improved by running the authorization function using Lambda@Edge closest to the viewer, reducing latency and response time to the viewer request. The load on your origin servers is also reduced by offloading CPU-intensive operations such as verification of JSON Web Token (JWT) signatures. Finally, there are security benefits such as filtering out unauthorized requests before they reach your origin infrastructure.

<https://aws.amazon.com/blogs/networking-and-content-delivery/authorizationedge-how-to-use-lambdaedge-and->

NEW QUESTION 191

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

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

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

HTTP 400 error (Bad request).

The response header also includes a status code element with a value of "Throttling" and a status message element with a value of "Rate exceeded "
Which combination of steps should the Solutions Architect take to resolve these issues? (Select THREE)

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

Answer: BEF

NEW QUESTION 195

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SAP-C01 Practice Exam Features:

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

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