

Amazon-Web-Services

Exam Questions SAP-C01

AWS Certified Solutions Architect- Professional



NEW QUESTION 1

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

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 2

A large company has many business units. Each business unit has multiple AWS accounts for different purposes. The CIO of the company sees that each business unit has data that would be useful to share with other parts of the company. In total, there are about 10 PB of data that needs to be shared with users in 1,000 AWS accounts. The data is proprietary, so some of it should only be available to users with specific job types. Some of the data is used for throughput of intensive workloads, such as simulations. The number of accounts changes frequently because of new initiatives, acquisitions, and divestitures.

A Solutions Architect has been asked to design a system that will allow for sharing data for use in AWS with all of the employees in the company.

Which approach will allow for secure data sharing in scalable way?

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

Answer: B

NEW QUESTION 3

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

- An application tier that contains 8 m4.xlarge instances
- A Classic Load Balancer
- Amazon S3 as a persistent data store

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

What should the Solution Architect recommend?

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

Answer: B

Explanation:

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

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

NEW QUESTION 4

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

- The website should be responsive.

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

How can the design requirements be met?

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

Answer: C

NEW QUESTION 5

A company currently uses Amazon EBS and Amazon RDS for storage purposes. The company intends to use a pilot light approach for disaster recovery in a different AWS Region. The company has an RTO of 6 hours and an RPO of 24 hours.

Which solution would achieve the requirements with MINIMAL cost?

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

Answer: A

Explanation:

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

NEW QUESTION 6

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 7

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 8

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

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

Answer: B

Explanation:

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

NEW QUESTION 9

As a part of building large applications in the AWS Cloud, the Solutions Architect is required to implement the perimeter security protection. Applications running on AWS have the following endpoints:

- Application Load Balancer
- Amazon API Gateway regional endpoint
- Elastic IP address-based EC2 instances.
- Amazon S3 hosted websites.
- Classic Load Balancer

The Solutions Architect must design a solution to protect all of the listed web front ends and provide the following security capabilities:

- DDoS protection
- SQL injection protection
- IP address whitelist/blacklist
- HTTP flood protection
- Bad bot scraper protection

How should the Solutions Architect design the solution?

- A. Deploy AWS WAF and AWS Shield Advanced on all web endpoints.
- B. Add AWS WAF rules to enforce the company's requirements.
- C. Deploy Amazon CloudFront in front of all the endpoints.
- D. The CloudFront distribution provides perimeter protection.
- E. Add AWS Lambda-based automation to provide additional security.
- F. Deploy Amazon CloudFront in front of all the endpoints.
- G. Deploy AWS WAF and AWS Shield Advanced.
- H. Add AWS WAF rules to enforce the company's requirements.
- I. Use AWS Lambda to automate and enhance the security posture.
- J. Secure the endpoints by using network ACLs and security groups and adding rules to enforce the company's requirements.
- K. Use AWS Lambda to automatically update the rules.

Answer: C

NEW QUESTION 10

A company CFO recently analyzed the company's AWS monthly bill and identified an opportunity to reduce the cost for AWS Elastic Beanstalk environments in use. The CFO has asked a Solutions Architect to design a highly available solution that will spin up an Elastic Beanstalk environment in the morning and terminate it at the end of the day.

The solution should be designed with minimal operational overhead and to minimize costs. It should also be able to handle the increased use of Elastic Beanstalk environments among different teams, and must provide a one-stop scheduler solution for all teams to keep the operational costs low.

What design will meet these requirements?

- A. Set up a Linux EC2 Micro instance.
- B. Configure an IAM role to allow the start and stop of the Elastic Beanstalk environment and attach it to the instance.
- C. Create scripts on the instance to start and stop the Elastic Beanstalk environment.
- D. Configure cron jobs on the instance to execute the scripts.

- E. Develop AWS Lambda functions to start and stop the Elastic Beanstalk environmen
- F. Configure a Lambda execution role granting Elastic Beanstalk environment start/stop permissions, and assign the role to the Lambda function
- G. Configure cron expression Amazon CloudWatch Events rules to trigger the Lambda functions.
- H. Develop an AWS Step Functions state machine with “wait” as its type to control the start and stop time. Use the activity task to start and stop the Elastic Beanstalk environmen
- I. Create a role for Step Functionsto allow it to start and stop the Elastic Beanstalk environmen
- J. Invoke Step Functions daily.
- K. Configure a time-based Auto Scaling grou
- L. In the morning, have the Auto Scaling group scale up an Amazon EC2 instance and put the Elastic Beanstalk environment start command in the EC2 instance user dat
- M. At the end of the day, scale down the instance number to 0 to terminate the EC2 instance.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/schedule-elastic-beanstalk-stop-restart/>

NEW QUESTION 10

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

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

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

Answer: DE

NEW QUESTION 11

A company is using AWS CloudFormation to deploy its infrastructure. The company is concerned that, if a production CloudFormation stack is deleted, important data stored in Amazon RDS databases or Amazon EBS volumes might also be deleted.

How can the company prevent users from accidentally deleting data in this way?

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

Answer: A

Explanation:

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

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

NEW QUESTION 14

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 18

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 21

A company is running an email application across multiple AWS Regions. The company uses Ohio (us-east-2) as the primary Region and Northern Virginia (us-east-1) as the Disaster Recovery (DR) Region. The data is continuously replicated from the primary Region to the DR Region by a single instance on the public subnet in both Regions. The replication messages between the Regions have a significant backlog during certain times of the day. The backlog clears on its own after a short time, but it affects the application's RPO.

Which of the following solutions should help remediate this performance problem? (Select TWO)

- A. Increase the size of the instances.
- B. Have the instance in the primary Region write the data to an Amazon SQS queue in the primary Region instead, and have the instance in the DR Region poll from this queue.
- C. Use multiple instances on the primary and DR Regions to send and receive the replication data.
- D. Change the DR Region to Oregon (us-west-2) instead of the current DR Region.
- E. Attach an additional elastic network interface to each of the instances in both Regions and set up load balancing between the network interfaces.

Answer: AC

NEW QUESTION 26

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

Which design meets these requirements?

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

Answer: B

NEW QUESTION 30

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 31

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 32

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

Which solution will meet the requirements at the LOWEST cost?

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

Answer: A

Explanation:

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

NEW QUESTION 34

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 39

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

How can a Solutions Architect achieve the isolation requirements?

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

Answer: A

NEW QUESTION 43

A company wants to manage the costs associated with a group of 20 applications that are critical, by migrating to AWS. The applications are a mix of Java and Node.js spread across different instance clusters. The company wants to minimize costs while standardizing by using a single deployment methodology. Most of the applications are part of month-end processing routines with a small number of concurrent users, but they are occasionally run at other times. Average application memory consumption is less than 1 GB, though some applications use as much as 2.5 GB of memory during peak processing. The most important application in the group is a billing report written in Java that accesses multiple data sources and often for several hours. Which is the MOST cost-effective solution?

- A. Deploy a separate AWS Lambda function for each applicatio
- B. Use AWS CloudTrail logs and Amazon CloudWatch alarms to verify completion of critical jobs.
- C. Deploy Amazon ECS containers on Amazon EC2 with Auto Scaling configured for memory utilization of 75%. Deploy an ECS task for each application being migrated with ECS task scalin
- D. Monitor services and hosts by using Amazon CloudWatch.
- E. Deploy AWS Elastic Beanstalk for each application with Auto Scaling to ensure that all requests have sufficient resource
- F. Monitor each AWS Elastic Beanstalk deployment with using CloudWatch alarms.
- G. Deploy a new amazon EC2 instance cluster that co-hosts all applications by using EC2 Auto Scaling and Application Load Balancer
- H. Scale cluster size based on a custom metric set on instance memory utilizatio
- I. Purchase 3-year Reserved instance reservations equal to the GroupMaxSize parameter of the Auto Scaling group.

Answer: C

NEW QUESTION 46

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 49

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

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

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

Answer: C

NEW QUESTION 50

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 51

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

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

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

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

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

Answer: D

NEW QUESTION 56

An advisory firm is creating a secure data analytics solution for its regulated financial services users. Users will upload their raw data to an Amazon S3 bucket, where they have PutObject permissions only. Data will be analyzed by applications running on an Amazon EMR cluster launched in a VPC. The firm requires that the environment be isolated from the internet. All data at rest must be encrypted using keys controlled by the firm. Which combination of actions should the Solutions Architect take to meet the user's security requirements? (Select TWO)

- A. Launch the Amazon EMR cluster in a private subnet configured to use an AWS KMS CMK for at-rest encryption. Configure a gateway VPC endpoint (for Amazon S3) and an interface VPC endpoint for AWS KMS.
- B. Launch the Amazon EMR cluster in a private subnet configured to use an AWS KMS CMK for at-rest encryption. Configure a gateway VPC endpoint for Amazon S3 and a NAT gateway to access AWS KMS.
- C. Launch the Amazon EMR cluster in a private subnet configured to use an AWS CloudHSM appliance for at-rest encryption. Configure a gateway VPC endpoint for Amazon S3 and an interface VPC endpoint for CloudHSM.
- D. Configure the S3 endpoint policies to permit access to the necessary data buckets only.
- E. Configure the S3 bucket policies to permit access using an aws:sourceVpce condition to match the S3 endpoint ID.

Answer: AC

NEW QUESTION 61

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 65

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

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

- A. Set up an AWS Storage Gateway, file gateway appliance on-premise
- B. Use the MAM solution to extract the videos from the current archive and push them into the file gateway
- C. Use the catalog of faces to build a collection in Amazon Rekognition
- D. Build an AWS Lambda function that invokes the Rekognition Javascript SDK to have Rekognition pull the video from the Amazon S3 files backing the file gateway, retrieve the required metadata, and push the metadata into the MAM solution.
- E. Set up an AWS Storage Gateway, tape gateway appliance on-premise
- F. Use the MAM solution to extract the videos from the current archive and push them into the tape gateway
- G. Use the catalog of faces to build a collection in Amazon Rekognition
- H. Build an AWS Lambda function that invokes the Rekognition Javascript SDK to have Amazon Rekognition process the video in the tape gateway, retrieve the required metadata, and push the metadata into the MAM solution.
- I. Configure a video ingestion stream by using Amazon Kinesis Video Stream
- J. Use the catalog of faces to build a collection in Amazon Rekognition
- K. Stream the videos from the MAM solution into Kinesis Video Stream
- L. Configure Amazon Rekognition to process the streamed video
- M. Then, use a stream consumer to retrieve the required metadata, and push the metadata into the MAM solution
- N. Configure the stream to store the videos in Amazon S3.
- O. Set up an Amazon EC2 instance that runs the OpenCV libraries
- P. Copy the videos, images, and face catalog from the on-premises library into an Amazon EBS volume mounted on this EC2 instance
- Q. Process the videos to retrieve the required metadata, and push the metadata into the MAM solution while also copying the video files to an Amazon S3 bucket.

Answer: C

Explanation:

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

NEW QUESTION 70

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 75

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

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

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

Answer: B

NEW QUESTION 80

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

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

Which design meets these requirements?

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

Answer: D

NEW QUESTION 82

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

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

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

Answer: B

Explanation:

https://aws.amazon.com/redshift/faqs/?nc1=h_ls 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 85

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 enable
T. Wait until the data is fully loaded and switch over the database connections to the new databas
. 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> I (DMS in transit encryption)
https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Security.html

NEW QUESTION 86

A company has a legacy application running on servers on premises. To increase the application's reliability, the company wants to gain actionable insights using application logs. A Solutions Architect has been given following requirements for the solution:

- Aggregate logs using AWS.
- Automate log analysis for errors.
- Notify the Operations team when errors go beyond a specified threshold. What solution meets the requirements?

A. Install Amazon Kinesis Agent on servers, send logs to Amazon Kinesis Data Streams and use Amazon Kinesis Data Analytics to identify errors, create an Amazon CloudWatch alarm to notify the Operations team of errors
B. Install an AWS X-Ray agent on servers, send logs to AWS Lambda and analyze them to identify errors, use Amazon CloudWatch Events to notify the Operations team of errors.
C. Install Logstash on servers, send logs to Amazon S3 and use Amazon Athena to identify errors, use sendmail to notify the Operations team of errors.
D. Install the Amazon CloudWatch agent on servers, send logs to Amazon CloudWatch Logs and use metric filters to identify errors, create a CloudWatch alarm to notify the Operations team of errors.

Answer: A

Explanation:

<https://docs.aws.amazon.com/kinesis-agent-windows/latest/userguide/what-is-kinesis-agent-windows.html> <https://medium.com/@khandelwal12nidhi/build-log-analytic-solution-on-aws-cc62a70057b2>

NEW QUESTION 89

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

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

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

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

Answer: CEF

NEW QUESTION 91

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 92

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

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

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



A single identity store is used to authenticate Developers across all companies. Which of the following approaches would meet these requirements? (Choose two.)

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

Answer: AD

NEW QUESTION 94

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 99

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 bucke
- D. Enable versioning on the Amazon S3 bucke
- E. Create Amazon S3 lifecycle policies to automatically migrate old versions of objects to Amazon S3 Standard 0 Infrequent Access, then Amazon Glacier, then delete backups after 1 year.
- F. Replace the local source code repository storage with a Storage Gateway stored volum
- G. Change the default snapshot frequency to 1 hou
- H. Use Amazon S3 lifecycle policies to archive snapshots to Amazon Glacier and remove old snapshots after 1 yea
- 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 volum
- K. Create a snapshot schedule to take hourly snapshot
- 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 104

A company has more than 100 AWS accounts, with one VPC per account, that need outbound HTTPS connectivity to the internet. The current design contains one NAT gateway per Availability Zone (AZ) in each VPC. To reduce costs and obtain information about outbound traffic, management has asked for a new architecture for internet access. Which solution will meet the current needs, and continue to grow as new accounts are provisioned, while reducing costs?

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

Answer: D

Explanation:

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

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

NEW QUESTION 107

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 110

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

Which workflow will meet these requirements in an automated manner?

- A. Add a Patch Group tag with a value of Windows Servers to all existing EC2 instances.
- B. Ensure that all Windows EC2 instances are assigned this tag.
- C. Associate the AWS-DefaultPatchBaseline to the Windows servers patch group.
- D. Define an AWS Systems Manager maintenance window, conduct patching within it, and associate it with the Windows Servers patch group.
- E. Register instances with the maintenance window using associated subnet ID.
- F. Assign the AWS-RunPatchBaseline document as a task within each maintenance window.
- G. Add a Patch Group tag with a value of Windows Servers to all existing EC2 instances.
- H. Ensure that all Windows EC2 instances are assigned this tag.
- I. Associate the AWS-WindowsPatchBaseline document as a task associated with the Windows Servers patch group.
- J. Create an Amazon CloudWatch Events rule configured to use a cron expression to schedule the execution of patching using the AWS Systems Manager run command.
- K. Create an AWS Systems Manager State Manager document to define commands to be executed during patch execution.
- L. Add a Patch Group tag with a value of either Windows Servers1 or Windows Server2 to all existing EC2 instances.
- M. Ensure that all Windows EC2 instances are assigned this tag.
- N. Associate the AWS-DefaultPatchBaseline with both Windows Servers patch groups.
- O. Define two non-overlapping AWS Systems Manager maintenance windows, conduct patching within them, and associate each with a different patch group.
- P. Register targets with specific maintenance windows using the Patch Group tag.
- Q. Assign the AWS-RunPatchBaseline document as a task within each maintenance window.
- R. Add a Patch Group tag with a value of either Windows Servers1 or Windows Server2 to all existing EC2 instances.
- S. Ensure that all Windows EC2 instances are assigned this tag.
- T. Associate the AWS-WindowsPatchBaseline with both Windows Servers patch groups.
- . Define two non-overlapping AWS Systems Manager maintenance windows, conduct patching within them, and associate each with a different patch group.
- . Assign the AWS-RunWindowsPatchBaseline document as a task within each maintenance window.
- . Create an AWS Systems Manager State Manager document to define commands to be executed during patch execution.

Answer: C

NEW QUESTION 113

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

Which solution is the MOST cost-effective at scale?

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

Answer: C

NEW QUESTION 115

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 118

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 121

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 125

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

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

Which of the following solutions is MOST cost-effective?

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

Answer: B

NEW QUESTION 127

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

steady 20% increase in utilization every month.

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

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

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

Answer: D

NEW QUESTION 131

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 133

A company is finalizing the architecture for its backup solution for applications running on AWS. All of the applications run on AWS and use at least two Availability Zones in each tier.

Company policy requires IT to durably store nightly backups for all its data in at least two locations: production and disaster recovery. The locations must be in different geographic regions. The company also needs the backup to be available to restore immediately at the production data center, and within 24 hours at the disaster recovery location. All backup processes must be fully automated.

What is the MOST cost-effective backup solution that will meet all requirements?

- A. Back up all the data to a large Amazon EBS volume attached to the backup media server in the production region
- B. Run automated scripts to snapshot these volumes nightly, and copy these snapshots to the disaster recovery region.
- C. Back up all the data to Amazon S3 in the disaster recovery region
- D. Use a lifecycle policy to move this data to Amazon Glacier in the production region immediately
- E. Only the data is replicated; remove the data from the S3 bucket in the disaster recovery region.
- F. Back up all the data to Amazon Glacier in the production region
- G. Set up cross-region replication of this data to Amazon Glacier in the disaster recovery region
- H. Set up a lifecycle policy to delete any data older than 60 days.
- I. Back up all the data to Amazon S3 in the production region
- J. Set up cross-region replication of this S3 bucket to another region and set up a lifecycle policy in the second region to immediately move this data to Amazon Glacier.

Answer: D

NEW QUESTION 135

A company is running a large application on-premises. Its technology stack consists of Microsoft .NET for the web server platform and Apache Cassandra for the database. The company wants to migrate the application to AWS to improve service reliability. The IT team also wants to reduce the time it spends on capacity management and maintenance of this infrastructure. The Development team is willing and available to make code changes to support the migration.

Which design is the LEAST complex to manage after the migration?

- A. Migrate the web servers to Amazon EC2 instances in an Auto Scaling group that is running .NET
- B. Migrate the existing Cassandra database to Amazon Aurora with multiple read replicas, and run both in a Multi-AZ mode.
- C. Migrate the web servers to an AWS Elastic Beanstalk environment that is running the .NET platform in a Multi-AZ Auto Scaling configuration

- D. Migrate the Cassandra database to Amazon EC2 instances that are running in a Multi-AZ configuration.
- E. Migrate the web servers to an AWS Elastic Beanstalk environment that is running the .NET platform in a Multi-AZ Auto Scaling configuration.
- F. Migrate the existing Cassandra database to Amazon DynamoDB.
- G. Migrate the web servers to Amazon EC2 instances in an Auto Scaling group that is running .NET.
- H. Migrate the existing Cassandra database to Amazon DynamoDB.

Answer: B

NEW QUESTION 137

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

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

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 138

A company has an application behind a load balancer with enough Amazon EC2 instances to satisfy peak demand. Scripts and third-party deployment solutions are used to configure EC2 instances when demand increases or an instance fails. The team must periodically evaluate the utilization of the instance types to ensure that the correct sizes are deployed.

How can this workload be optimized to meet these requirements?

- A. Use CloudFormation to create AWS CloudFormation stacks from the current resource.
- B. Deploy that stack by using AWS CloudFormation in the same region.
- C. Use Amazon CloudWatch alarms to send notifications about underutilized resources to provide cost-savings suggestions.
- D. Create an Auto Scaling group to scale the instances, and use AWS CodeDeploy to perform the configuration.
- E. Change from a load balancer to an Application Load Balancer.
- F. Purchase a third-party product that provides suggestions for cost savings on AWS resources.
- G. Deploy the application by using AWS Elastic Beanstalk with default option.
- H. Register for an AWS Support Developer plan.
- I. Review the instance usage for the application by using Amazon CloudWatch, and identify less expensive instances that can handle the load.
- J. Hold monthly meetings to review new instance types and determine whether Reserved instances should be purchased.
- K. Deploy the application as a Docker image by using Amazon ECS.
- L. Set up Amazon EC2 Auto Scaling and Amazon ECS scaling.
- M. Register for AWS Business Support and use Trusted Advisor checks to provide suggestions on cost savings.

Answer: D

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 to 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 143

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

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

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

- A. Create a policy that allows creation of project-related resources only.
- B. Create roles with required service permissions, which are assumable by the services.
- C. Create a policy that allows creation of all project-related resources, including roles that allow access only to specified resources.
- D. Create roles with the required service permissions, which are assumable by the service.

- E. Have the interns create and use a bastion host to create the project resources in the project subnet only.
- F. Create a policy that allows creation of project-related resources only.
- G. Require the interns to raise a request for roles to be created with the Security team.
- H. The interns will provide the requirements for the permissions to be set in the role.

Answer: A

NEW QUESTION 145

A company wants to move a web application to AWS. The application stores session information locally on each web server, which will make auto scaling difficult. As part of the migration, the application will be rewritten to decouple the session data from the web servers. The company requires low latency, scalability, and availability.

Which service will meet the requirements for storing the session information in the MOST cost-effective way?

- A. Amazon ElastiCache with the Memcached engine
- B. Amazon S3
- C. Amazon RDS MySQL
- D. Amazon ElastiCache with the Redis engine

Answer: D

Explanation:

<https://aws.amazon.com/caching/session-management/> <https://aws.amazon.com/elasticache/redis-vs-memcached/>

NEW QUESTION 149

A company has a data center that must be migrated to AWS as quickly as possible. The data center has a 500 Mbps AWS Direct Connect link and a separate, fully available 1 Gbps ISP connection. A Solutions Architect must transfer 20 TB of data from the data center to an Amazon S3 bucket.

What is the FASTEST way to transfer the data?

- A. Upload the data to the S3 bucket using the existing DX link.
- B. Send the data to AWS using the AWS Import/Export service.
- C. Upload the data using an 80 TB AWS Snowball device.
- D. Upload the data to the S3 bucket using S3 Transfer Acceleration.

Answer: D

Explanation:

<https://aws.amazon.com/s3/faqs/>

NEW QUESTION 152

A company is having issues with a newly deployed serverless infrastructure that uses Amazon API Gateway, Amazon Lambda, and Amazon DynamoDB.

In a steady state, the application performs as expected. However, during peak load, tens of thousands of simultaneous invocations are needed and user requests fail multiple times before succeeding. The company has checked the logs for each component, focusing specifically on Amazon CloudWatch Logs for Lambda.

There are no errors logged by the services or applications.

What might cause this problem?

- A. Lambda has very little memory assigned, which causes the function to fail at peak load.
- B. Lambda is in a subnet that uses a NAT gateway to reach out to the internet, and the function instance does not have sufficient Amazon EC2 resources in the VPC to scale with the load.
- C. The throttle limit set on API Gateway is very low during peak load, the additional requests are not making their way through to Lambda.
- D. DynamoDB is set up in an auto scaling mode.
- E. During peak load, DynamoDB adjusts capacity and through successfully.

Answer: A

NEW QUESTION 154

A large global company wants to migrate a stateless mission-critical application to AWS. The application is based on IBM WebSphere (application and integration middleware), IBM MQ (messaging middleware), and IBM DB2 (database software) on a z/OS operating system.

How should the Solutions Architect migrate the application to AWS?

- A. Re-host WebSphere-based applications on Amazon EC2 behind a load balancer with Auto Scaling. Re-platform the IBM MQ to an Amazon EC2-based MQ.
- B. Re-platform the z/OS-based DB2 to Amazon RDS DB2.
- C. Re-host WebSphere-based applications on Amazon EC2 behind a load balancer with Auto Scaling. Re-platform the IBM MQ to an Amazon MQ.
- D. Re-platform z/OS-based DB2 to Amazon EC2-based DB2.
- E. Orchestrate and deploy the application by using AWS Elastic Beanstalk.
- F. Re-platform the IBM MQ to Amazon SQS.
- G. Re-platform z/OS-based DB2 to Amazon RDS DB2.
- H. Use the AWS Server Migration Service to migrate the IBM WebSphere and IBM DB2 to an Amazon EC2-based solution.
- I. Re-platform the IBM MQ to an Amazon MQ.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/database/aws-database-migration-service-and-aws-schema-conversion-tool-now->
<https://aws.amazon.com/quickstart/architecture/ibm-mq/>

NEW QUESTION 159

A public retail web application uses an Application Load Balancer (ALB) in front of Amazon EC2 instances running across multiple Availability Zones (AZs) in a Region backed by an Amazon RDS MySQL Multi-AZ deployment. Target group health checks are configured to use HTTP and pointed at the product catalog

page. Auto Scaling is configured to maintain the web fleet size based on the ALB health check.

Recently, the application experienced an outage. Auto Scaling continuously replaced the instances during the outage. A subsequent investigation determined that the web server metrics were within the normal range, but the database tier was experiencing high load, resulting in severely elevated query response times.

Which of the following changes together would remediate these issues while improving monitoring capabilities for the availability and functionality of the entire application stack for future growth? (Select TWO.)

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

Answer: BE

NEW QUESTION 164

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

Which of the following would speed up this process?

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

Answer: A

Explanation:

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

NEW QUESTION 166

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

Which option meets the requirements and MINIMIZES costs?

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

Answer: A

NEW QUESTION 167

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 170

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 175

A company prefers to limit running Amazon EC2 instances to those that were launched from AMIs pre-approved by the Information Security department. The Development team has an agile continuous integration and deployment process that cannot be stalled by the solution.

Which method enforces the required controls with the LEAST impact on the development process? (Choose two.)

- A. Use IAM policies to restrict the ability of users or other automated entities to launch EC2 instances based on a specific set of pre-approved AMIs, such as those tagged in a specific way by Information Security.
- B. Use regular scans within Amazon Inspector with a custom assessment template to determine if the EC2 instance that the Amazon Inspector Agent is running on is based upon a pre-approved AM
- C. If it is not, shut down the instance and inform information Security by email that this occurred.
- D. Only allow launching of EC2 instances using a centralized DevOps team, which is given work packages via notifications from an internal ticketing syste
- E. Users make requests for resources using this ticketing tool, which has manual information security approval steps to ensure that EC2 instances are only launched from approved AMIs.
- F. Use AWS Config rules to spot any launches of EC2 instances based on non-approved AMIs, trigger an AWS Lambda function to automatically terminate the instance, and publish a message to an Amazon SNS topic to inform Information Security that this occurred.
- G. Use a scheduled AWS Lambda function to scan through the list of running instances within the virtual private cloud (VPC) and determine if any of these are based on unapproved AMI
- H. Publish a message to an SNS topic to inform Information Security that this occurred and then shut down the instance.

Answer: AD

Explanation:

https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config_develop-rules_getting-started.html

NEW QUESTION 178

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

Currently the company uses a combination of Amazon CloudWatch Events and AWS Lambda to create the DNS record. The solution worked well in testing with small clusters, but in production with clusters containing thousands of instances the company sees the following error in the Lambda logs:

HTTP 400 error (Bad request).

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

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

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

Answer: BEF

NEW QUESTION 181

.....

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](#)