

Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional



NEW QUESTION 1

- (Exam Topic 1)

A company has a three-tier application running on AWS with a web server, an application server, and an Amazon RDS MySQL DB instance. A solutions architect is designing a disaster recovery (DR) solution with an RPO of 5 minutes.

Which solution will meet the company's requirements?

- A. Configure AWS Backup to perform cross-Region backups of all servers every 5 minute
- B. Reprovision the three tiers in the DR Region from the backups using AWS CloudFormation in the event of a disaster.
- C. Maintain another running copy of the web and application server stack in the DR Region using AWS CloudFormation drill detectio
- D. Configure cross-Region snapshots of the DB instance to the DR Region every 5 minute
- E. In the event of a disaster, restore the DB instance using the snapshot in the DR Region.
- F. Use Amazon EC2 Image Builder to create and copy AMIs of the web and application server to both the primary and DR Region
- G. Create a cross-Region read replica of the DB instance in the DR Regio
- H. In the event of a disaster, promote the read replica to become the master and reprovision the servers with AWS CloudFormation using the AMIs.
- I. Create AMIs of the web and application servers in the DR Regio
- J. Use scheduled AWS Glue jobs to synchronize the DB instance with another DB instance in the DR Regio
- K. In the event of a disaster, switch to the DB instance in the DR Region and reprovision the servers with AWS CloudFormation using the AMIs.

Answer: C

Explanation:

deploying a brand new RDS instance will take >30 minutes. You will use EC2 Image builder to put the AMIs into the new region, but not use image builder to LAUNCH them.

NEW QUESTION 2

- (Exam Topic 1)

A company has a project that is launching Amazon EC2 instances that are larger than required. The project's account cannot be part of the company's organization in AWS Organizations due to policy restrictions to keep this activity outside of corporate IT. The company wants to allow only the launch of t3.small EC2 instances by developers in the project's account. These EC2 instances must be restricted to the us-east-2 Region.

What should a solutions architect do to meet these requirements?

- A. Create a new developer accoun
- B. Move all EC2 instances, users, and assets into us-east-2. Add the account to the company's organization in AWS Organization
- C. Enforce a tagging policy that denotes Region affinity.
- D. Create an SCP that denies the launch of all EC2 instances except I3.small EC2 instances in us-east-2. Attach the SCP to the project's account.
- E. Create and purchase a t3.small EC2 Reserved Instance for each developer in us-east-2. Assign each developer a specific EC2 instance with their name as the tag.
- F. Create an IAM policy than allows the launch of only t3.small EC2 instances in us-east-2. Attach the policy to the roles and groups that the developers use in the project's account.

Answer: D

NEW QUESTION 3

- (Exam Topic 1)

A company wants to control its cost of Amazon Athena usage The company has allocated a specific monthly budget for Athena usage A solutions architect must design a solution that will prevent the company from exceeding the budgeted amount

Which solution will moot these requirements?

- A. Use AWS Budget
- B. Create an alarm (or when the cost of Athena usage reaches the budgeted amount for the mont
- C. Configure AWS Budgets actions to deactivate Athena until the end of the month.
- D. Use Cost Explorer to create an alert for when the cost of Athena usage reaches the budgeted amount for the mont
- E. Configure Cost Explorer to publish notifications to an Amazon Simple Notification Service (Amazon SNS) topic.
- F. Use AWS Trusted Advisor to track the cost of Athena usag
- G. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule to deactivate Athena until the end of the month whenever the cost reaches the budgeted amount for the month
- H. Use Athena workgroups to set a limit on the amount of data that can be scanne
- I. Set a limit that is appropriate for the monthly budget and the current pricing for Athena.

Answer: D

NEW QUESTION 4

- (Exam Topic 1)

A team collects and routes behavioral data for an entire company The company runs a Multi-AZ VPC environment with public subnets, private subnets, and in internet gateway Each public subnet also contains a NAT gateway Most of the company's applications read from and write to Amazon Kinesis Data Streams. Most of the workloads am in private subnets.

A solutions architect must review the infrastructure The solutions architect needs to reduce costs and maintain the function of the applications The solutions architect uses Cost Explorer and notices that the cost in the EC2-Other category is consistently high A further review shows that NatGateway-Bytes charges are increasing the cost in the EC2-Other category.

What should the solutions architect do to meet these requirements?

- A. Enable VPC Flow Log
- B. Use Amazon Athena to analyze the logs for traffic that can be remove
- C. Ensure that security groups are Mocking traffic that is responsible for high costs.
- D. Add an interface VPC endpoint for Kinesis Data Streams to the VP
- E. Ensure that applications have the correct IAM permissions to use the interface VPC endpoint.
- F. Enable VPC Flow Logs and Amazon Detective Review Detective findings for traffic that is not related to Kinesis Data Streams Configure security groups to block that traffic

- G. Add an interface VPC endpoint for Kinesis Data Streams to the VP
- H. Ensure that the VPC endpoint policy allows traffic from the applications.

Answer: D

Explanation:

<https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-access.html> <https://aws.amazon.com/premiumsupport/knowledge-center/vpc-reduce-nat-gateway-transfer-costs/>

VPC endpoint policies enable you to control access by either attaching a policy to a VPC endpoint or by using additional fields in a policy that is attached to an IAM user, group, or role to restrict access to only occur via the specified VPC endpoint

NEW QUESTION 5

- (Exam Topic 1)

A company needs to store and process image data that will be uploaded from mobile devices using a custom mobile app. Usage peaks between 8 AM and 5 PM on weekdays, with thousands of uploads per minute. The app is rarely used at any other time A user is notified when image processing is complete. Which combination of actions should a solutions architect take to ensure image processing can scale to handle the load? (Select THREE.)

- A. Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon MQ queue.
- B. Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon Simple Queue Service (Amazon SQS) standard queue.
- C. Invoke an AWS Lambda function to perform image processing when a message is available in the queue.
- D. Invoke an S3 Batch Operations job to perform image processing when a message is available in the queue.
- E. Send a push notification to the mobile app by using Amazon Simple Notification Service (Amazon SNS) when processing is complete.
- F. Send a push notification to the mobile app by using Amazon Simple Email Service (Amazon SES) when processing is complete.

Answer: BCE

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/batch-ops-basics.html>

NEW QUESTION 6

- (Exam Topic 1)

A company is running a traditional web application on Amazon EC2 instances. The company needs to refactor the application as microservices that run on containers. Separate versions of the application exist in two distinct environments: production and testing. Load for the application is variable, but the minimum load and the maximum load are known. A solutions architect needs to design the updated application with a serverless architecture that minimizes operational complexity.

Which solution will meet these requirements MOST cost-effectively?

- A. Upload the container images to AWS Lambda as function
- B. Configure a concurrency limit for the associated Lambda functions to handle the expected peak loa
- C. Configure two separate Lambda integrations within Amazon API Gateway: one for production and one for testing.
- D. Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Container Service (Amazon ECS) clusters with the Fargate launch type to handle the expected loa
- E. Deploy tasks from the ECR image
- F. Configure two separate Application Load Balancers to direct traffic to the ECS clusters.
- G. Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Kubernetes Service (Amazon EKS) clusters with the Fargate launch type to handle the expected loa
- H. Deploy tasks from the ECR image
- I. Configure two separate Application Load Balancers to direct traffic to the EKS clusters.
- J. Upload the container images to AWS Elastic Beanstalk
- K. In Elastic Beanstalk, create separate environments and deployments for production and testin
- L. Configure two separate Application Load Balancers to direct traffic to the Elastic Beanstalk deployments.

Answer: D

Explanation:

minimizes operational + microservices that run on containers = AWS Elastic Beanstalk

NEW QUESTION 7

- (Exam Topic 1)

A company has application services that have been containerized and deployed on multiple Amazon EC2 instances with public IPs. An Apache Kafka cluster has been deployed to the EC2 instances. A PostgreSQL database has been migrated to Amazon RDS for PostgreSQL. The company expects a significant increase of orders on its platform when a new version of its flagship product is released.

What changes to the current architecture will reduce operational overhead and support the product release?

- A. Create an EC2 Auto Scaling group behind an Application Load Balance
- B. Create additional read replicas for the DB instanc
- C. Create Amazon Kinesis data streams and configure the application services lo use the data stream
- D. Store and serve static content directly from Amazon S3.
- E. Create an EC2 Auto Scaling group behind an Application Load Balance
- F. Deploy the DB instance in Multi-AZ mode and enable storage auto scalin
- G. Create Amazon Kinesis data streams and configure the application services to use the data stream
- H. Store and serve static content directly from Amazon S3.
- I. Deploy the application on a Kubernetes cluster created on the EC2 instances behind an Application Load Balance
- J. Deploy the DB instance in Multi-AZ mode and enable storage auto scalin
- K. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluste
- L. Store static content in Amazon S3 behind an Amazon CloudFront distribution.
- M. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balance
- N. Create additional read replicas for the DB instanc
- O. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluste

P. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

Answer: D

Explanation:

Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balancer. Create additional read replicas for the DB instance. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

NEW QUESTION 8

- (Exam Topic 1)

A large company in Europe plans to migrate its applications to the AWS Cloud. The company uses multiple AWS accounts for various business groups. A data privacy law requires the company to restrict developers' access to AWS European Regions only.

What should the solutions architect do to meet this requirement with the LEAST amount of management overhead?

- A. Create IAM users and IAM groups in each account
- B. Create IAM policies to limit access to non-European Regions Attach the IAM policies to the IAM groups
- C. Enable AWS Organizations, attach the AWS accounts, and create OUs for European Regions and non-European Region
- D. Create SCPs to limit access to non-European Regions and attach the policies to the OUs.
- E. Set up AWS Single Sign-On and attach AWS account
- F. Create permission sets with policies to restrict access to non-European Regions Create IAM users and IAM groups in each account.
- G. Enable AWS Organizations, attach the AWS accounts, and create OUs for European Regions and non-European Region
- H. Create permission sets with policies to restrict access to non-European Region
- I. Create IAM users and IAM groups in the primary account.

Answer: B

Explanation:

"This policy uses the Deny effect to deny access to all requests for operations that don't target one of the two approved regions (eu-central-1 and eu-west-1)."

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_general.htm

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements_condition.html

NEW QUESTION 9

- (Exam Topic 1)

A company is building an image service on the web that will allow users to upload and search random photos. At peak usage, up to 10,000 users worldwide will upload their images. The service will then overlay text on the uploaded images, which will then be published on the company website.

Which design should a solutions architect implement?

- A. Store the uploaded images in Amazon Elastic File System (Amazon EFS). Send application log information about each image to Amazon CloudWatch Log
- B. Create a fleet of Amazon EC2 instances that use CloudWatch Logs to determine which images need to be processed
- C. Place processed images in another directory in Amazon EFS
- D. Enable Amazon CloudFront and configure the origin to be the one of the EC2 instances in the fleet.
- E. Store the uploaded images in an Amazon S3 bucket and configure an S3 bucket event notification to send a message to Amazon Simple Notification Service (Amazon SNS). Create a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) to pull messages from Amazon SNS to process the images and place them in Amazon Elastic File System (Amazon EFS). Use Amazon CloudWatch metrics for the SNS message volume to scale out EC2 instance
- F. Enable Amazon CloudFront and configure the origin to be the ALB in front of the EC2 instances.
- G. Store the uploaded images in an Amazon S3 bucket and configure an S3 bucket event notification to send a message to the Amazon Simple Queue Service (Amazon SQS) queue
- H. Create a fleet of Amazon EC2 instances to pull messages from the SQS queue to process the images and place them in another S3 bucket
- I. Use Amazon CloudWatch metrics for queue depth to scale out EC2 instance
- J. Enable Amazon CloudFront and configure the origin to be the S3 bucket that contains the processed images.
- K. Store the uploaded images on a shared Amazon Elastic Block Store (Amazon EBS) volume mounted to a fleet of Amazon EC2 Spot instance
- L. Create an Amazon DynamoDB table that contains information about each uploaded image and whether it has been processed
- M. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to scale out EC2 instance
- N. Enable Amazon CloudFront and configure the origin to reference an Elastic Load Balancer in front of the fleet of EC2 instances.

Answer: C

NEW QUESTION 10

- (Exam Topic 1)

A developer reports receiving an Error 403: Access Denied message when they try to download an object from an Amazon S3 bucket. The S3 bucket is accessed using an S3 endpoint inside a VPC, and is encrypted with an AWS KMS key. A solutions architect has verified that the developer is assuming the correct IAM role in the account that allows the object to be downloaded. The S3 bucket policy and the NACL are also valid.

Which additional step should the solutions architect take to troubleshoot this issue?

- A. Ensure that blocking all public access has not been enabled in the S3 bucket.
- B. Verify that the IAM role has permission to decrypt the referenced KMS key.
- C. Verify that the IAM role has the correct trust relationship configured.
- D. Check that local firewall rules are not preventing access to the S3 endpoint.

Answer: B

NEW QUESTION 10

- (Exam Topic 1)

A large company with hundreds of AWS accounts has a newly established centralized internal process for purchasing new or modifying existing Reserved Instances. This process requires all business units that want to purchase or modify Reserved Instances to submit requests to a dedicated team for procurement or execution. Previously, business units would directly purchase or modify Reserved Instances in their own respective AWS accounts autonomously.

Which combination of steps should be taken to proactively enforce the new process in the MOST secure way possible? (Select TWO.)

- A. Ensure all AWS accounts are part of an AWS Organizations structure operating in all features mode.
- B. Use AWS Contig lo report on the attachment of an IAM policy that denies access to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions.
- C. In each AWS account, create an IAM policy with a DENY rule to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions.
- D. Create an SCP that contains a deny rule to the ec2:PurchaseReservedInstancesOffering and ec2: Modify Reserved Instances action
- E. Attach the SCP to each organizational unit (OU) of the AWS Organizations structure.
- F. Ensure that all AWS accounts are part of an AWS Organizations structure operating in consolidated billing features mode.

Answer: AD

Explanation:

https://docs.aws.amazon.com/organizations/latest/APIReference/API_EnableAllFeatures.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp-strategies.html

NEW QUESTION 13

- (Exam Topic 1)

A company has a new application that needs to run on five Amazon EC2 instances in a single AWS Region. The application requires high-throughput, low-latency network connections between all of the EC2 instances where the application will run. There is no requirement for the application to be fault tolerant. Which solution will meet these requirements?

- A. Launch five new EC2 instances into a cluster placement grou
- B. Ensure that the EC2 instance type supports enhanced networking.
- C. Launch five new EC2 instances into an Auto Scaling group in the same Availability Zon
- D. Attach an extra elastic network interface to each EC2 instance.
- E. Launch five new EC2 instances into a partition placement grou
- F. Ensure that the EC2 instance type supports enhanced networking.
- G. Launch five new EC2 instances into a spread placement grou
- H. Attach an extra elastic network interface to each EC2 instance.

Answer: A

Explanation:

When you launch EC2 instances in a cluster they benefit from performance and low latency. No redundancy though as per the question
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>.

NEW QUESTION 16

- (Exam Topic 1)

A solutions architect is designing a network for a new cloud deployment. Each account will need autonomy to modify route tables and make changes. Centralized and controlled egress internet connectivity is also needed. The cloud footprint is expected to grow to thousands of AWS accounts. Which architecture will meet these requirements?

- A. A centralized transit VPC with a VPN connection to a standalone VPC in each accoun
- B. Outbound internet traffic will be controlled by firewall appliances.
- C. A centralized shared VPC with a subnet for each accoun
- D. Outbound internet traffic will controlled through a fleet of proxy servers.
- E. A shared services VPC to host central assets to include a fleet of firewalls with a route to the internet.Each spoke VPC will peer to the central VPC.
- F. A shared transit gateway to which each VPC will be attache
- G. Outbound internet access will route through a fleet of VPN-attached firewalls.

Answer: D

Explanation:

<https://docs.aws.amazon.com/whitepapers/latest/building-scalable-secure-multi-vpc-network-infrastructure/centr>
<https://docs.aws.amazon.com/whitepapers/latest/building-scalable-secure-multi-vpc-network-infrastructure/centr>
AWS Transit Gateway helps you design and implement networks at scale by acting as a cloud router. As your network grows, the complexity of managing incremental connections can slow you down. AWS Transit Gateway connects VPCs and on-premises networks through a central hub. This simplifies your network and puts an end to complex peering relationships -- each new connection is only made once.

NEW QUESTION 18

- (Exam Topic 1)

A company with global offices has a single 1 Gbps AWS Direct Connect connection to a single AWS Region. The company's on-premises network uses the connection to communicate with the company's resources in the AWS Cloud. The connection has a single private virtual interface that connects to a single VPC. A solutions architect must implement a solution that adds a redundant Direct Connect connection in the same Region. The solution also must provide connectivity to other Regions through the same pair of Direct Connect connections as the company expands into other Regions. Which solution meets these requirements?

- A. Provision a Direct Connect gatewa
- B. Delete the existing private virtual interface from the existing connectio
- C. Create the second Direct Connect connectio
- D. Create a new private virtual interlace on each connection, and connect both private virtual interfaces to the Direct Connect gatewa
- E. Connect the Direct Connect gateway to the single VPC.
- F. Keep the existing private virtual interfac
- G. Create the second Direct Connect connectio
- H. Create a new private virtual interface on the new connection, and connect the new private virtual interface to the single VPC.
- I. Keep the existing private virtual interfac
- J. Create the second Direct Connect connectio
- K. Create a new public virtual interface on the new connection, and connect the new public virtual interface to the single VPC.
- L. Provision a transit gatewa
- M. Delete the existing private virtual interface from the existing connection.Create the second Direct Connect connectio
- N. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the transit gatewa

O. Associate the transit gateway with the single VPC.

Answer: A

Explanation:

A Direct Connect gateway is a globally available resource. You can create the Direct Connect gateway in any Region and access it from all other Regions. The following describe scenarios where you can use a Direct Connect gateway.

<https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html>

NEW QUESTION 22

- (Exam Topic 1)

An ecommerce website running on AWS uses an Amazon RDS for MySQL DB instance with General Purpose SSD storage. The developers chose an appropriate instance type based on demand, and configured 100 GB of storage with a sufficient amount of free space.

The website was running smoothly for a few weeks until a marketing campaign launched. On the second day of the campaign, users reported long wait times and time outs. Amazon CloudWatch metrics indicated that both reads and writes to the DB instance were experiencing long response times. The CloudWatch metrics show 40% to 50% CPU and memory utilization, and sufficient free storage space is still available. The application server logs show no evidence of database connectivity issues.

What could be the root cause of the issue with the marketing campaign?

- A. It exhausted the I/O credit balance due to provisioning low disk storage during the setup phase.
- B. It caused the data in the tables to change frequently, requiring indexes to be rebuilt to optimize queries.
- C. It exhausted the maximum number of allowed connections to the database instance.
- D. It exhausted the network bandwidth available to the RDS for MySQL DB instance.

Answer: A

Explanation:

"When using General Purpose SSD storage, your DB instance receives an initial I/O credit balance of 5.4 million I/O credits. This initial credit balance is enough to sustain a burst performance of 3,000 IOPS for 30 minutes."

<https://aws.amazon.com/blogs/database/how-to-use-cloudwatch-metrics-to-decide-between-general-purpose-or>

NEW QUESTION 24

- (Exam Topic 1)

A company is running an application on Amazon EC2 instances in three environments; development, testing, and production. The company uses AMIs to deploy the EC2 instances. The company builds the AMIs by using custom deployment scripts and infrastructure orchestration tools for each release in each environment. The company is receiving errors in its deployment process. Errors appear during operating system package downloads and during application code installation from a third-party Git hosting service. The company needs deployments to become more reliable across all environments.

Which combination of steps will meet these requirements? (Select THREE).

- A. Mirror the application code to an AWS CodeCommit Git repositior
- B. Use the repository to build EC2 AMIs.
- C. Produce multiple EC2 AMI
- D. one for each environment, for each release.
- E. Produce one EC2 AMI for each release for use across all environments.
- F. Mirror the application code to a third-party Git repository that uses Amazon S3 storag
- G. Use the repository for deployment.
- H. Replace the custom scripts and tools with AWS CodeBuil
- I. Update the infrastructure deployment process to use EC2 Image Builder.

Answer: ACE

NEW QUESTION 25

- (Exam Topic 1)

A company has an internal application running on AWS that is used to track and process shipments in the company's warehouse. Currently, after the system receives an order, it emails the staff the information needed to ship a package. Once the package is shipped, the staff replies to the email and the order is marked as shipped.

The company wants to stop using email in the application and move to a serverless application model. Which architecture solution meets these requirements?

- A. Use AWS Batch to configure the different tasks required lo ship a packag
- B. Have AWS Batch trigger an AWS Lambda function that creates and prints a shipping labe
- C. Once that label is scanne
- D. as it leaves the warehouse, have another Lambda function move the process to the next step in the AWS Batch job.B.
- E. When a new order is created, store the order information in Amazon SQ
- F. Have AWS Lambda check the queue every 5 minutes and process any needed wor
- G. When an order needs to be shipped, have Lambda print the label in the warehous
- H. Once the label has been scanned, as it leaves the warehouse, have an Amazon EC2 instance update Amazon SOS.
- I. Update the application to store new order information in Amazon DynamoD
- J. When a new order is created, trigger an AWS Step Functions workflow, mark the orders as "in progress," and print a package label to the warehous
- K. Once the label has been scanned and fulfilled, the application will trigger an AWS Lambda function that will mark the order as shipped and complete the workflow.
- L. Store new order information in Amazon EF
- M. Have instances pull the new information from the NFS and send that information to printers in the warehous
- N. Once the label has been scanned, as it leaves the warehouse, have Amazon API Gateway call the instances to remove the order information from Amazon EFS.

Answer: C

NEW QUESTION 30

- (Exam Topic 1)

A media company uses Amazon DynamoDB to store metadata for its catalog of movies that are available to stream. Each media item contains user-facing content that includes a description of the media, a list of search tags, and similar data. In addition, media items include a list of Amazon S3 key names that relate to movie files. The company stores these movie files in a single S3 bucket that has versioning enabled. The company uses Amazon CloudFront to serve these movie files. The company has 100,000 media items, and each media item can have many different S3 objects that represent different encodings of the same media. S3 objects that belong to the same media item are grouped together under the same key prefix, which is a random unique ID. Because of an expiring contract with a media provider, the company must remove 2,000 media items. The company must completely delete all DynamoDB keys and movie files on Amazon S3 that are related to these media items within 36 hours. The company must ensure that the content cannot be recovered. Which combination of actions will meet these requirements? (Select TWO.)

- A. Configure the DynamoDB table with a TTL field
- B. Create and invoke an AWS Lambda function to perform a conditional update. Set the TTL field to the time of the contract's expiration on every affected media item.
- C. Configure an S3 Lifecycle object expiration rule that is based on the contract's expiration date
- D. Write a script to perform a conditional delete on all the affected DynamoDB records
- E. Temporarily suspend versioning on the S3 bucket
- F. Create and invoke an AWS Lambda function that deletes affected objects. Reactivate versioning when the operation is complete
- G. Write a script to delete objects from Amazon S3. Specify in each request a NoncurrentVersionExpiration property with a NoncurrentDays attribute set to 0.

Answer: CE

NEW QUESTION 35

- (Exam Topic 1)

A travel company built a web application that uses Amazon Simple Email Service (Amazon SES) to send email notifications to users. The company needs to enable logging to help troubleshoot email delivery issues. The company also needs the ability to do searches that are based on recipient, subject, and time sent. Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

- A. Create an Amazon SES configuration set with Amazon Kinesis Data Firehose as the destination
- B. Choose to send logs to an Amazon S3 bucket.
- C. Enable AWS CloudTrail logging
- D. Specify an Amazon S3 bucket as the destination for the logs.
- E. Use Amazon Athena to query the logs in the Amazon S3 bucket for recipient, subject, and time sent.
- F. Create an Amazon CloudWatch log group
- G. Configure Amazon SES to send logs to the log group
- H. Use Amazon Athena to query the logs in Amazon CloudWatch for recipient, subject, and time sent.

Answer: AC

Explanation:

<https://docs.aws.amazon.com/ses/latest/dg/event-publishing-retrieving-firehose.html>

To enable you to track your email sending at a granular level, you can set up Amazon SES to publish email sending events to Amazon CloudWatch, Amazon Kinesis Data Firehose, or Amazon Simple Notification Service based on characteristics that you define.

<https://docs.aws.amazon.com/ses/latest/dg/monitor-using-event-publishing.html>

<https://aws.amazon.com/getting-started/hands-on/build-serverless-real-time-data-processing-app-lambda-kinesis>

NEW QUESTION 36

- (Exam Topic 1)

A company hosts a web application that runs on a group of Amazon EC2 instances that are behind an Application Load Balancer (ALB) in a VPC. The company wants to analyze the network payloads to reverse-engineer a sophisticated attack of the application. Which approach should the company take to achieve this goal?

- A. Enable VPC Flow Log
- B. Store the flow logs in an Amazon S3 bucket for analysis.
- C. Enable Traffic Mirroring on the network interface of the EC2 instance
- D. Send the mirrored traffic to a target for storage and analysis.
- E. Create an AWS WAF web ACL
- F. and associate it with the ALB
- G. Configure AWS WAF logging.
- H. Enable logging for the ALB
- I. Store the logs in an Amazon S3 bucket for analysis.

Answer: A

NEW QUESTION 39

- (Exam Topic 1)

A company needs to implement a patching process for its servers. The on-premises servers and Amazon EC2 instances use a variety of tools to perform patching. Management requires a single report showing the patch status of all the servers and instances. Which set of actions should a solutions architect take to meet these requirements?

- A. Use AWS Systems Manager to manage patches on the on-premises servers and EC2 instances
- B. Use Systems Manager to generate patch compliance reports.
- C. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instances
- D. Use Amazon QuickSight integration with OpsWorks to generate patch compliance reports.
- E. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to apply patches by scheduling an AWS Systems Manager patch remediation job
- F. Use Amazon Inspector to generate patch compliance reports.
- G. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instances
- H. Use AWS X-Ray to post the patch status to AWS Systems Manager OpsCenter to generate patch compliance reports.

Answer: A

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-patch.html>

NEW QUESTION 44

- (Exam Topic 1)

A company has registered 10 new domain names. The company uses the domains for online marketing. The company needs a solution that will redirect online visitors to a specific URL for each domain. All domains and target URLs are defined in a JSON document. All DNS records are managed by Amazon Route 53. A solutions architect must implement a redirect service that accepts HTTP and HTTPS requests.

Which combination of steps should the solutions architect take to meet these requirements with the LEAST amount of operational effort? (Choose three.)

- A. Create a dynamic webpage that runs on an Amazon EC2 instance
- B. Configure the webpage to use the JSON document in combination with the event message to look up and respond with a redirect URL.
- C. Create an Application Load Balancer that includes HTTP and HTTPS listeners.
- D. Create an AWS Lambda function that uses the JSON document in combination with the event message to look up and respond with a redirect URL.
- E. Use an Amazon API Gateway API with a custom domain to publish an AWS Lambda function.
- F. Create an Amazon CloudFront distribution
- G. Deploy a Lambda@Edge function.
- H. Create an SSL certificate by using AWS Certificate Manager (ACM). Include the domains as Subject Alternative Names.

Answer: CEF

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-edge-how-it-works-tutorial.html>

NEW QUESTION 45

- (Exam Topic 1)

A large payroll company recently merged with a small staffing company. The unified company now has multiple business units, each with its own existing AWS account.

A solutions architect must ensure that the company can centrally manage the billing and access policies for all the AWS accounts. The solutions architect configures AWS Organizations by sending an invitation to all member accounts of the company from a centralized management account.

What should the solutions architect do next to meet these requirements?

- A. Create the OrganizationAccountAccess IAM group in each member account
- B. Include the necessary IAM roles for each administrator.
- C. Create the OrganizationAccountAccessPolicy IAM policy in each member account
- D. Connect the member accounts to the management account by using cross-account access.
- E. Create the OrganizationAccountAccessRole IAM role in each member account
- F. Grant permission to the management account to assume the IAM role.
- G. Create the OrganizationAccountAccessRole IAM role in the management account Attach the Administrator Access AWS managed policy to the IAM role
- H. Assign the IAM role to the administrators in each member account.

Answer: C

NEW QUESTION 46

- (Exam Topic 1)

A company is planning on hosting its ecommerce platform on AWS using a multi-tier web application designed for a NoSQL database. The company plans to use the us-west-2 Region as its primary Region. The company wants to ensure that copies of the application and data are available in a second Region, us-west-1, for disaster recovery. The company wants to keep the time to fail over as low as possible. Failing back to the primary Region should be possible without administrative interaction after the primary service is restored.

Which design should the solutions architect use?

- A. Use AWS CloudFormation StackSets to create the stacks in both Regions with Auto Scaling groups for the web and application tier
- B. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication
- C. Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outage
- D. Use Amazon DynamoDB global tables for the database tier.
- E. Use AWS CloudFormation StackSets to create the stacks in both Regions with Auto Scaling groups for the web and application tier
- F. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication
- G. Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outage
- H. Deploy an Amazon Aurora global database for the database tier.
- I. Use AWS Service Catalog to deploy the web and application servers in both Regions
- J. Asynchronously replicate static content between the two Regions using Amazon S3 cross-Region replication
- K. Use Amazon Route 53 health checks to identify a primary Region failure and update the public DNS entry listing to the secondary Region in the event of an outage
- L. Use Amazon RDS for MySQL with cross-Region replication for the database tier.
- M. Use AWS CloudFormation StackSets to create the stacks in both Regions using Auto Scaling groups for the web and application tier
- N. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication
- O. Use Amazon CloudFront with static files in Amazon S3, and multi-Region origins for the front-end web tier
- P. Use Amazon DynamoDB tables in each Region with scheduled backups to Amazon S3.

Answer: A

NEW QUESTION 47

- (Exam Topic 1)

A company is running an application distributed over several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The security team requires that all application access attempts be made available for analysis. Information about the client IP address, connection type, and user agent must be included.

Which solution will meet these requirements?

- A. Enable EC2 detailed monitoring, and include network logs. Send all logs through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon

ES) cluster that the security team uses for analysis.

- B. Enable VPC Flow Logs for all EC2 instance network interfaces Publish VPC Flow Logs to an Amazon S3 bucket Have the security team use Amazon Athena to query and analyze the logs
- C. Enable access logs for the Application Load Balancer, and publish the logs to an Amazon S3 bucket Have the security team use Amazon Athena to query and analyze the logs
- D. Enable Traffic Mirroring and specify all EC2 instance network interfaces as the source
- E. Send all traffic information through Amazon Kinesis Data Firehose to an Amazon Elastic search Service (Amazon ES) cluster that the security team uses for analysis.

Answer: C

Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html>

NEW QUESTION 49

- (Exam Topic 1)

A company uses a service to collect metadata from applications that the company hosts on premises. Consumer devices such as TVs and internet radios access the applications. Many older devices do not support certain HTTP headers and exhibit errors when these headers are present in responses. The company has configured an on-premises load balancer to remove the unsupported headers from responses sent to older devices, which the company identified by the User-Agent headers.

The company wants to migrate the service to AWS, adopt serverless technologies, and retain the ability to support the older devices. The company has already migrated the applications into a set of AWS Lambda functions.

Which solution will meet these requirements?

- A. Create an Amazon CloudFront distribution for the metadata service
- B. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the ALB
- C. Configure the ALB to invoke the correct Lambda function for each type of request
- D. Create a CloudFront function to remove the problematic headers based on the value of the User-Agent header.
- E. Create an Amazon API Gateway REST API for the metadata service
- F. Configure API Gateway to invoke the correct Lambda function for each type of request
- G. Modify the default gateway responses to remove the problematic headers based on the value of the User-Agent header.
- H. Create an Amazon API Gateway HTTP API for the metadata service
- I. Configure API Gateway to invoke the correct Lambda function for each type of request
- J. Create a response mapping template to remove the problematic headers based on the value of the User-Agent header
- K. Associate the response data mapping with the HTTP API.
- L. Create an Amazon CloudFront distribution for the metadata service
- M. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the ALB
- N. Configure the ALB to invoke the correct Lambda function for each type of request
- O. Create a Lambda@Edge function that will remove the problematic headers in response to viewer requests based on the value of the User-Agent header.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html>

NEW QUESTION 52

- (Exam Topic 1)

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

Explanation:

Ephemeral ports are not covered in the syllabus so be careful that you don't confuse day to day best practice with what is required for the exam. Link to an explanation on Ephemeral ports here. <https://acloud.guru/forums/aws-certified-solutions-architect-associate/discussion/-KUbCwo4IXefMI7janaK/netw>

NEW QUESTION 54

- (Exam Topic 1)

A company provides a centralized Amazon EC2 application hosted in a single shared VPC. The centralized application must be accessible from client applications running in the VPCs of other business units. The centralized application front end is configured with a Network Load Balancer (NLB) for scalability.

Up to 10 business unit VPCs will need to be connected to the shared VPC. Some of the business unit VPC CIDR blocks overlap with the shared VPC, and some overlap with each other. Network connectivity to the centralized application in the shared VPC should be allowed from authorized business unit VPCs only.

Which network configuration should a solutions architect use to provide connectivity from the client applications in the business unit VPCs to the centralized application in the shared VPC?

- A. Create an AWS Transit Gateway
- B. Attach the shared VPC and the authorized business unit VPCs to the transit gateway
- C. Create a single transit gateway route table and associate it with all of the attached VPCs
- D. Allow automatic propagation of routes from the attachments into the route table
- E. Configure VPC routing tables to send traffic to the transit gateway.

- F. Create a VPC endpoint service using the centralized application NLB and enable (he option to require endpoint acceptanc
- G. Create a VPC endpoint in each of the business unit VPCs using the service name of the endpoint servic
- H. Accept authorized endpoint requests from the endpoint service console.
- I. Create a VPC peering connection from each business unit VPC to lthe shared VP
- J. Accept the VPC peering connections from the shared VPC consol
- K. Configure VPC routing tables to send traffic to theVPC peering connection.
- L. Configure a virtual private gateway for the shared VPC and create customer gateways for each of the authorized business unit VPC
- M. Establish a Sile-to-Site VPN connection from the business unit VPCs to the shared VP
- N. Configure VPC routing tables to send traffic to the VPN connection.

Answer: B

Explanation:

Amazon Transit Gateway doesn't support routing between Amazon VPCs with overlapping CIDRs. If you attach a new Amazon VPC that has a CIDR which overlaps with an already attached Amazon VPC, Amazon Transit Gateway will not propagate the new Amazon VPC route into the Amazon Transit Gateway route table.

<https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-target-groups.html#client-ip-pre>

NEW QUESTION 59

- (Exam Topic 1)

A company hosts a large on-premises MySQL database at its main office that supports an issue tracking system used by employees around the world. The company already uses AWS for some workloads and has created an Amazon Route 53 entry for the database endpoint that points to the on-premises database. Management is concerned about the database being a single point of failure and wants a solutions architect to migrate the database to AWS without any data loss or downtime.

Which set of actions should the solutions architect implement?

- A. Create an Amazon Aurora DB cluste
- B. Use AWS Database Migration Service (AWS DMS) to do a full load from the on-premises database lo Auror
- C. Update the Route 53 entry for the database to point to the Aurora cluster endpoint
- D. and shut down the on-premises database.
- E. During nonbusiness hours, shut down the on-premises database and create a backu
- F. Restore this backup to an Amazon Aurora DB cluste
- G. When the restoration is complete, update the Route 53 entry for the database to point to the Aurora cluster endpoint, and shut down the on-premises database.
- H. Create an Amazon Aurora DB cluste
- I. Use AWS Database Migration Service (AWS DMS) to do a full load with continuous replication from the on-premises database to Auror
- J. When the migration is complete, update the Route 53 entry for the database to point to the Aurora cluster endpoint, and shut down the on-premises database.
- K. Create a backup of the database and restore it to an Amazon Aurora multi-master cluste
- L. This Aurora cluster will be in a master-master replication configuration with the on-premises databas
- M. Update the Route 53 entry for the database to point to the Aurora cluster endpoint
- N. and shut down the on-premises database.

Answer: C

Explanation:

"Around the world" eliminates possibility for the maintenance window at night. The other difference is ability to leverage continuous replication in MySQL to Aurora case.

NEW QUESTION 61

- (Exam Topic 1)

A company runs an e-commerce platform with front-end and e-commerce tiers. Both tiers run on LAMP stacks with the front-end instances running behind a load balancing appliance that has a virtual offering on AWS Current*/, the operations team uses SSH to log in to the instances to maintain patches and address other concerns. The platform has recently been the target of multiple attacks, including.

- A DDoS attack.
- An SQL injection attack
- Several successful dictionary attacks on SSH accounts on the web servers

The company wants to improve the security of the e-commerce platform by migrating to AWS. The company's solutions architects have decided to use the following approach;

- Code review the existing application and fix any SQL injection issues.
- Migrate the web application to AWS and leverage the latest AWS Linux AMI to address initial security patching.
- Install AWS Systems Manager to manage patching and allow the system administrators to run commands on all instances, as needed.

What additional steps will address all of the identified attack types while providing high availability and minimizing risk?

- A. Enable SSH access to the Amazon EC2 instances using a security group that limits access to specific IP
- B. Migrate on-premises MySQL to Amazon RDS Multi-AZ Install the third-party load balancer from the AWS Marketplace and migrate the existing rules to the load balancer's AWS instances Enable AWS Shield Standard for DDoS protection
- C. Disable SSH access to the Amazon EC2 instance
- D. Migrate on-premises MySQL to Amazon RDS Multi-AZ Leverage an Elastic Load Balancer to spread the load and enable AWS Shield Advanced for protectio
- E. Add an Amazon CloudFront distribution in front of the website Enable AWS WAF on the distribution to manage the rules.
- F. Enable SSH access to the Amazon EC2 instances through a bastion host secured by limiting access to specific IP addresse
- G. Migrate on-premises MySQL to a self-managed EC2 instanc
- H. Leverage an AWS Elastic Load Balancer to spread the load, and enable AWS Shield Standard for DDoS protection Add anAmazon CloudFront distribution in front of the website.
- I. Disable SSH access to the EC2 instance
- J. Migrate on-premises MySQL to Amazon RDS Single-A
- K. Leverage an AWS Elastic Load Balancer to spread the load Add an Amazon CloudFront distribution in front of the website Enable AWS WAF on the distribution to manage the rules.

Answer: B

NEW QUESTION 66

- (Exam Topic 1)

A company is running an Apache Hadoop cluster on Amazon EC2 instances. The Hadoop cluster stores approximately 100 TB of data for weekly operational reports and allows occasional access for data scientists to retrieve data. The company needs to reduce the cost and operational complexity for storing and serving this data.

Which solution meets these requirements in the MOST cost-effective manner?

- A. Move the Hadoop cluster from EC2 instances to Amazon EM
- B. Allow data access patterns to remain the same.
- C. Write a script that resizes the EC2 instances to a smaller instance type during downtime and resizes the instances to a larger instance type before the reports are created.
- D. Move the data to Amazon S3 and use Amazon Athena to query the data for report
- E. Allow the data scientists to access the data directly in Amazon S3.
- F. Migrate the data to Amazon DynamoDB and modify the reports to fetch data from DynamoD
- G. Allow the data scientists to access the data directly in DynamoDB.

Answer: C

Explanation:

"The company needs to reduce the cost and operational complexity for storing and serving this data. Which solution meets these requirements in the MOST cost-effective manner?" EMR storage is ephemeral. The company has 100TB that need to persist, they would have to use EMRFS to backup to S3 anyway.

<https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-plan-storage.html>

100TB

EBS - 8.109\$ S3 - 2.355\$

You have saved 5.752\$

This amount can be used for Athen. BTW. we don't know indexes, amount of data that is scanned. What we know is that it will be: "occasional access for data scientists to retrieve data"

NEW QUESTION 70

- (Exam Topic 1)

A company has implemented an ordering system using an event-driven architecture. During initial testing, the system stopped processing orders. Further analysis revealed that one order message in an Amazon Simple Queue Service (Amazon SQS) standard queue was causing an error on the backend and blocking all subsequent order messages. The visibility timeout of the queue is set to 30 seconds, and the backend processing timeout is set to 10 seconds. A solutions architect needs to analyze faulty order messages and ensure that the system continues to process subsequent messages.

Which step should the solutions architect take to meet these requirements?

- A. Increase the backend processing timeout to 30 seconds to match the visibility timeout
- B. Reduce the visibility timeout of the queue to automatically remove the faulty message
- C. Configure a new SQS FIFO queue as a dead-letter queue to isolate the faulty messages
- D. Configure a new SQS standard queue as a dead-letter queue to isolate the faulty messages.

Answer: D

NEW QUESTION 75

- (Exam Topic 1)

A company has a multi-tier web application that runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB and the Auto Scaling group are replicated in a backup AWS Region. The minimum value and the maximum value for the Auto Scaling group are set to zero. An Amazon RDS Multi-AZ DB instance stores the application's data. The DB instance has a read replica in the backup Region. The application presents an endpoint to end users by using an Amazon Route 53 record.

The company needs to reduce its RTO to less than 15 minutes by giving the application the ability to automatically fail over to the backup Region. The company does not have a large enough budget for an active-active strategy.

What should a solutions architect recommend to meet these requirements?

- A. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALB
- B. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
- C. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Region
- D. Configure the CloudWatch alarm to invoke the Lambda function.
- E. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
- F. Configure Route 53 with a health check that monitors the web application and sends an Amazon Simple Notification Service (Amazon SNS) notification to the Lambda function when the health check status is unhealthy
- G. Update the application's Route 53 record with a failover policy that routes traffic to the ALB in the backup Region when a health check failure occurs.
- H. Configure the Auto Scaling group in the backup Region to have the same values as the Auto Scaling group in the primary Region
- I. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALB
- J. Remove the read replica
- K. Replace the read replica with a standalone RDS DB instance
- L. Configure Cross-Region Replication between the RDS DB instances by using snapshots and Amazon S3.
- M. Configure an endpoint in AWS Global Accelerator with the two ALBs as equal weighted targets
- N. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
- O. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Region
- P. Configure the CloudWatch alarm to invoke the Lambda function.

Answer: B

Explanation:

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

NEW QUESTION 76

- (Exam Topic 1)

A solutions architect is designing the data storage and retrieval architecture for a new application that a company will be launching soon. The application is designed to ingest millions of small records per minute from devices all around the world. Each record is less than 4 KB in size and needs to be stored in a durable

location where it can be retrieved with low latency. The data is ephemeral and the company is required to store the data for 120 days only, after which the data can be deleted.

The solutions architect calculates that, during the course of a year, the storage requirements would be about 10-15 TB.

Which storage strategy is the MOST cost-effective and meets the design requirements?

- A. Design the application to store each incoming record as a single .csv file in an Amazon S3 bucket to allow for indexed retrieval
- B. Configure a lifecycle policy to delete data older than 120 days.
- C. Design the application to store each incoming record in an Amazon DynamoDB table properly configured for the scale
- D. Configure the DynamoDB Time to Live (TTL) feature to delete records older than 120 days.
- E. Design the application to store each incoming record in a single table in an Amazon RDS MySQL database
- F. Run a nightly cron job that executes a query to delete any records older than 120 days.
- G. Design the application to batch incoming records before writing them to an Amazon S3 bucket
- H. Update the metadata for the object to contain the list of records in the batch and use the Amazon S3 metadata search feature to retrieve the data
- I. Configure a lifecycle policy to delete the data after 120 days.

Answer: B

Explanation:

DynamoDB with TTL, cheaper for sustained throughput of small items + suited for fast retrievals. S3 cheaper for storage only, much higher costs with writes. RDS not designed for this use case.

NEW QUESTION 78

- (Exam Topic 1)

A solutions architect has an operational workload deployed on Amazon EC2 instances in an Auto Scaling group. The VPC architecture spans two Availability Zones (AZ) with a subnet in each that the Auto Scaling group is targeting. The VPC is connected to an on-premises environment and connectivity cannot be interrupted. The maximum size of the Auto Scaling group is 20 instances in service. The VPC IPv4 addressing is as follows:

VPC CIDR: 10.0.0.0/23

AZ1 subnet CIDR: 10.0.0.0/24 AZ2 subnet CIDR: 10.0.1.0/24

Since deployment, a third AZ has become available in the Region. The solutions architect wants to adopt the new AZ without adding additional IPv4 address space and without service downtime.

Which solution will meet these requirements?

- A. Update the Auto Scaling group to use the AZ2 subnet only
- B. Delete and re-create the AZ1 subnet using half the previous address space
- C. Adjust the Auto Scaling group to also use the new AZ1 subnet
- D. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only
- E. Remove the current AZ2 subnet
- F. Create a new AZ2 subnet using the second half of the address space from the original AZ1 subnet
- G. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.
- H. Terminate the EC2 instances in the AZ1 subnet
- I. Delete and re-create the AZ1 subnet using half the address space
- J. Update the Auto Scaling group to use this new subnet
- K. Repeat this for the second AZ
- L. Define a new subnet in AZ3, then update the Auto Scaling group to target all three new subnets.
- M. Create a new VPC with the same IPv4 address space and define three subnets, with one for each AZ
- N. Update the existing Auto Scaling group to target the new subnets in the new VPC.
- O. Update the Auto Scaling group to use the AZ2 subnet only
- P. Update the AZ1 subnet to have half the previous address space
- Q. Adjust the Auto Scaling group to also use the AZ1 subnet again
- R. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only
- S. Update the current AZ2 subnet and assign the second half of the address space from the original AZ1 subnet
- T. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.

Answer: A

Explanation:

https://aws.amazon.com/premiumsupport/knowledge-center/vpc-ip-address-range/?nc1=h_ls

It's not possible to modify the IP address range of an existing virtual private cloud (VPC) or subnet. You must delete the VPC or subnet, and then create a new VPC or subnet with your preferred CIDR block.

NEW QUESTION 83

- (Exam Topic 1)

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 data sharing 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. The policy should allow the accounts in the partnership read access to the bucket
- F. Enable Requester Pays on the bucket
- G. Have the organizations use their AWS credentials when accessing the data.
- H. Ensure that all organizations in the partnership have AWS account
- I. 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. Periodically sync the data from the institute's account to the other organization
- J. Have the organizations use their AWS credentials when accessing the data using their accounts
- K. Ensure that all organizations in the partnership have AWS account
- L. 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

- M. Enable Requester Pays on the bucket
- N. Have the organizations assume and use that read role when accessing the data.

Answer: B

Explanation:

In general, bucket owners pay for all Amazon S3 storage and data transfer costs associated with their bucket. A bucket owner, however, can configure a bucket to be a Requester Pays bucket. With Requester Pays buckets, the requester instead of the bucket owner pays the cost of the request and the data download from the bucket. The bucket owner always pays the cost of storing data. If you enable Requester Pays on a bucket, anonymous access to that bucket is not allowed.
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/RequesterPaysExamples.html>

NEW QUESTION 87

- (Exam Topic 1)

A company wants to deploy an AWS WAF solution to manage AWS WAF rules across multiple AWS accounts. The accounts are managed under different OUs in AWS Organizations.

Administrators must be able to add or remove accounts or OUs from managed AWS WAF rule sets as needed. Administrators also must have the ability to automatically update and remediate noncompliant AWS WAF rules in all accounts.

Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Use AWS Firewall Manager to manage AWS WAF rules across accounts in the organization.
- B. Use an AWS Systems Manager Parameter Store parameter to store account numbers and OUs to manage. Update the parameter as needed to add or remove accounts or OUs. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to identify any changes to the parameter and to invoke an AWS Lambda function to update the security policy in the Firewall Manager administrative account.
- C. Deploy an organization-wide AWS Config rule that requires all resources in the selected OUs to associate the AWS WAF rule.
- D. Deploy automated remediation actions by using AWS Lambda to fix noncompliant resources. Deploy AWS WAF rules by using an AWS CloudFormation stack set to target the same OUs where the AWS Config rule is applied.
- E. Create AWS WAF rules in the management account of the organization. Use AWS Lambda environment variables to store account numbers and OUs to manage. Update environment variables as needed to add or remove accounts or OUs. Create cross-account IAM roles in member accounts. Assume the roles by using AWS Security Token Service (AWS STS) in the Lambda function to create and update AWS WAF rules in the member accounts.
- F. Use AWS Control Tower to manage AWS WAF rules across accounts in the organization. Use AWS Key Management Service (AWS KMS) to store account numbers and OUs to manage. Update AWS KMS as needed to add or remove accounts or OUs. Create IAM users in member accounts. Allow AWS Control Tower in the management account to use the access key and secret access key to create and update AWS WAF rules in the member accounts.

Answer: D

NEW QUESTION 90

- (Exam Topic 1)

A company plans to migrate to AWS. A solutions architect uses AWS Application Discovery Service over the fleet and discovers that there is an Oracle data warehouse and several PostgreSQL databases. Which combination of migration patterns will reduce licensing costs and operational overhead? (Select TWO.)

- A. Lift and shift the Oracle data warehouse to Amazon EC2 using AWS DMS.
- B. Migrate the Oracle data warehouse to Amazon Redshift using AWS SCT and AWS QMS.
- C. Lift and shift the PostgreSQL databases to Amazon EC2 using AWS DMS.
- D. Migrate the PostgreSQL databases to Amazon RDS for PostgreSQL using AWS DMS.
- E. Migrate the Oracle data warehouse to an Amazon EMR managed cluster using AWS DMS.

Answer: BD

Explanation:

<https://aws.amazon.com/getting-started/hands-on/migrate-oracle-to-amazon-redshift/> <https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/migrate-an-on-premises-postgresql-database>

NEW QUESTION 95

- (Exam Topic 1)

A company is deploying a new cluster for big data analytics on AWS. The cluster will run across many Linux Amazon EC2 instances that are spread across multiple Availability Zones.

All of the nodes in the cluster must have read and write access to common underlying file storage. The file storage must be highly available, must be resilient, must be compatible with the Portable Operating System Interface (POSIX), and must accommodate high levels of throughput.

Which storage solution will meet these requirements?

- A. Provision an AWS Storage Gateway file gateway NFS file share that is attached to an Amazon S3 bucket.
- B. Mount the NFS file share on each EC2 instance in the cluster.
- C. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses General Purpose performance mode.
- D. Mount the EFS file system on each EC2 instance in the cluster.
- E. Provision a new Amazon Elastic Block Store (Amazon EBS) volume that uses the io2 volume type. Attach the EBS volume to all of the EC2 instances in the cluster.
- F. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses Max I/O performance mode.
- G. Mount the EFS file system on each EC2 instance in the cluster.

Answer: D

NEW QUESTION 100

- (Exam Topic 1)

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

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

- A. Pre-warming Elastic Load Balancers, using a bigger instance type, changing all Amazon EBS volumes to GP2 volumes.
- B. Performing a one-time migration of the database cluster to Amazon RD
- C. and creating several additional read replicas to handle the load during end of month
- D. Using Amazon CloudWatch with AWS Lambda to change the type
- E. size, or IOPS of Amazon EBS volumes in the cluster based on a specific CloudWatch metric
- F. 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

Explanation:

In this scenario, the Amazon EC2 instances are in an Auto Scaling group already which means that the database read operations is the possible bottleneck especially during the month-end wherein the reports are generated. This can be solved by creating RDS read replicas.

NEW QUESTION 101

- (Exam Topic 1)

A company is running a data-intensive application on AWS. The application runs on a cluster of hundreds of Amazon EC2 instances. A shared file system also runs on several EC2 instances that store 200 TB of data. The application reads and modifies the data on the shared file system and generates a report. The job runs once monthly, reads a subset of the files from the shared file system, and takes about 72 hours to complete. The compute instances scale in an Auto Scaling group, but the instances that host the shared file system run continuously. The compute and storage instances are all in the same AWS Region.

A solutions architect needs to reduce costs by replacing the shared file system instances. The file system must provide high performance access to the needed data for the duration of the 72-hour run.

Which solution will provide the LARGEST overall cost reduction while meeting these requirements?

- A. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Intelligent-Tiering storage class
- B. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using lazy loadin
- C. Use the new file system as the shared storage for the duration of the job
- D. Delete the file system when the job is complete.
- E. Migrate the data from the existing shared file system to a large Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enable
- F. Attach the EBS volume to each of the instances by using a user data script in the Auto Scaling group launch template
- G. Use the EBS volume as the shared storage for the duration of the job
- H. Detach the EBS volume when the job is complete.
- I. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Standard storage class
- J. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using batch loadin
- K. Use the new file system as the shared storage for the duration of the job
- L. Delete the file system when the job is complete.
- M. Migrate the data from the existing shared file system to an Amazon S3 bucket
- N. Before the job runs each month, use AWS Storage Gateway to create a file gateway with the data from Amazon S3. Use the file gateway as the shared storage for the job
- O. Delete the file gateway when the job is complete.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/storage/new-enhancements-for-moving-data-between-amazon-fsx-for-lustre-and>

NEW QUESTION 105

- (Exam Topic 1)

A company has a policy that all Amazon EC2 instances that are running a database must exist within the same subnets in a shared VPC. Administrators must follow security compliance requirements and are not allowed to directly log in to the shared account. All company accounts are members of the same organization in AWS Organizations. The number of accounts will rapidly increase as the company grows.

A solutions architect uses AWS Resource Access Manager to create a resource share in the shared account. What is the MOST operationally efficient configuration to meet these requirements?

- A. Add the VPC to the resource share
- B. Add the account IDs as principals
- C. Add all subnets within the VPC to the resource share
- D. Add the account IDs as principals
- E. Add all subnets within the VPC to the resource share
- F. Add the organization as a principal.
- G. Add the VPC to the resource share
- H. Add the organization as a principal

Answer: C

Explanation:

<https://docs.aws.amazon.com/ram/latest/userguide/getting-started-sharing.html#getting-started-sharing-create> To restrict resource sharing to only principals in your organization, choose Allow sharing with principals in your organization only.

<https://docs.aws.amazon.com/ram/latest/userguide/ram-ug.pdf>

NEW QUESTION 110

- (Exam Topic 1)

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

Which steps should the solutions architect take to meet these requirements?

- A. Create an AWS CloudHSM cluster
- B. Create a new CMK in AWS KMS using AWS_CloudHSM as the source (or the key material and an origin of AWS_CLOUDHSM)
- C. Enable automatic key rotation on the CMK with a duration of 1 year

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

Answer: C

Explanation:

<https://aws.amazon.com/blogs/security/how-to-byok-bring-your-own-key-to-aws-kms-for-less-than-15-00-a-year>
<https://docs.aws.amazon.com/kms/latest/developerguide/importing-keys-create-cmk.html>

NEW QUESTION 111

- (Exam Topic 1)

A company needs to run a software package that has a license that must be run on the same physical host for the duration of its use. The software package is only going to be used for 90 days. The company requires patching and restarting of all instances every 30 days. How can these requirements be met using AWS?

- A. Run a dedicated instance with auto-placement disabled.
- B. Run the instance on a dedicated host with Host Affinity set to Host.
- C. Run an On-Demand Instance with a Reserved Instance to ensure consistent placement.
- D. Run the instance on a licensed host with termination set for 90 days.

Answer: B

Explanation:

Host Affinity is configured at the instance level. It establishes a launch relationship between an instance and a Dedicated Host. (This sets which host the instance can run on) Auto-placement allows you to manage whether instances that you launch are launched onto a specific host, or onto any available host that has matching configurations. Auto-placement must be configured at the host level. (This sets which instance the host can run.) When affinity is set to Host, an instance launched onto a specific host always restarts on the same host if stopped. This applies to both targeted and untargeted launches.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/how-dedicated-hosts-work.html>

When affinity is set to Off, and you stop and restart the instance, it can be restarted on any available host. However, it tries to launch back onto the last Dedicated Host on which it ran (on a best-effort basis).

NEW QUESTION 112

- (Exam Topic 1)

A multimedia company needs to deliver its video-on-demand (VOD) content to its subscribers in a cost-effective way. The video files range in size from 1-15 GB and are typically viewed frequently for the first 6 months after creation, and then access decreases considerably. The company requires all video files to remain immediately available for subscribers. There are now roughly 30,000 files, and the company anticipates doubling that number over time.

What is the MOST cost-effective solution for delivering the company's VOD content?

- A. Store the video files in an Amazon S3 bucket using S3 Intelligent-Tiering
- B. Use Amazon CloudFront to deliver the content with the S3 bucket as the origin.
- C. Use AWS Elemental MediaConvert and store the adaptive bitrate video files in Amazon S3. Configure an AWS Elemental MediaPackage endpoint to deliver the content from Amazon S3.
- D. Store the video files in Amazon Elastic File System (Amazon EFS) Standard
- E. Enable EFS lifecycle management to move the video files to EFS Infrequent Access after 6 months
- F. Create an Amazon EC2 Auto Scaling group behind an Elastic Load Balancer to deliver the content from Amazon EFS.
- G. Store the video files in Amazon S3 Standard
- H. Create S3 Lifecycle rules to move the video files to S3 Standard-Infrequent Access (S3 Standard-IA) after 6 months and to S3 Glacier Deep Archive after 1 year
- I. Use Amazon CloudFront to deliver the content with the S3 bucket as the origin.

Answer: A

Explanation:

<https://d1.awsstatic.com/whitepapers/amazon-cloudfront-for-media.pdf> <https://aws.amazon.com/solutions/implementations/video-on-demand-on-aws/>

NEW QUESTION 115

- (Exam Topic 1)

A company has developed a single-page web application in JavaScript. The source code is stored in a single Amazon S3 bucket in the us-east-1 Region. The company serves the web application to a global user base through Amazon CloudFront.

The company wants to experiment with two versions of the website without informing application users. Each version of the website will reside in its own S3 bucket. The company wants to determine which version is most successful in marketing a new product.

The solution must send application users that are based in Europe to the new website design. The solution must send application users that are based in the United States to the current website design. However, some exceptions exist. The company needs to be able to redirect specific users to the new website design, regardless of the users' location.

Which solution meets these requirements?

- A. Configure two CloudFront distributions
- B. Configure a geolocation routing policy in Amazon Route 53 to route traffic to the appropriate CloudFront endpoint based on the location of clients.
- C. Configure a single CloudFront distribution

- D. Create a behavior with different paths for each version of the site
- E. Configure Lambda@Edge on the default path to generate redirects and send the client to the correct version of the website.
- F. Configure a single CloudFront distribution
- G. Configure an alternate domain name on the distribution. Configure two behaviors to route users to the different S3 origins based on the domain name that the client uses in the HTTP request.
- H. Configure a single CloudFront distribution with Lambda@Edge
- I. Use Lambda@Edge to send user requests to different origins based on request attributes.

Answer: A

NEW QUESTION 120

- (Exam Topic 1)

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

- * 1. The data must be highly durable and available.
- * 2. The data must always be encrypted at rest and in transit.
- * 3. The encryption key must be managed by the company and rotated periodically.

Which of the following solutions should the solutions architect recommend?

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

Answer: B

Explanation:

Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.

NEW QUESTION 125

- (Exam Topic 1)

A company is migrating applications from on-premises to the AWS Cloud. These applications power the company's internal web forms. These web forms collect data for specific events several times each quarter. The web forms use simple SQL statements to save the data to a local relational database.

Data collection occurs for each event, and the on-premises servers are idle most of the time. The company needs to minimize the amount of idle infrastructure that supports the web forms.

Which solution will meet these requirements?

- A. Use Amazon EC2 Image Builder to create AMIs for the legacy server
- B. Use the AMIs to provision EC2 instances to recreate the applications in the AWS Cloud
- C. Place an Application Load Balancer (ALB) in front of the EC2 instance
- D. Use Amazon Route 53 to point the DNS names of the web forms to the ALB.
- E. Create one Amazon DynamoDB table to store data for all the data input. Use the application form name as the table key to distinguish data items
- F. Create an Amazon Kinesis data stream to receive the data input and store the input in DynamoDB
- G. Use Amazon Route 53 to point the DNS names of the web forms to the Kinesis data stream's endpoint.
- H. Create Docker images for each server of the legacy web form application
- I. Create an Amazon Elastic Container Service (Amazon ECS) cluster on AWS Fargate
- J. Place an Application Load Balancer in front of the ECS cluster
- K. Use Fargate task storage to store the web form data.
- L. Provision an Amazon Aurora Serverless cluster
- M. Build multiple schemas for each web form's data storage
- N. Use Amazon API Gateway and an AWS Lambda function to recreate the data input form
- O. Use Amazon Route 53 to point the DNS names of the web forms to their corresponding API Gateway endpoint.

Answer: D

Explanation:

Provision an Amazon Aurora Serverless cluster. Build multiple schemas for each web form's data storage. Use Amazon API Gateway and an AWS Lambda function to recreate the data input forms. Use Amazon Route 53 to point the DNS names of the web forms to their corresponding API Gateway endpoint.

NEW QUESTION 128

- (Exam Topic 1)

A company is building a hybrid solution between its existing on-premises systems and a new backend in AWS. The company has a management application to monitor the state of its current IT infrastructure and automate responses to issues. The company wants to incorporate the status of its consumed AWS services into the application. The application uses an HTTPS endpoint to receive updates.

Which approach meets these requirements with the LEAST amount of operational overhead?

- A. Configure AWS Systems Manager OpsCenter to ingest operational events from the on-premises systems. Retire the on-premises management application and adopt OpsCenter as the hub
- B. Configure Amazon EventBridge (Amazon CloudWatch Events) to detect and react to changes for AWS Health events from the AWS Personal Health Dashboard. Configure the EventBridge (CloudWatch Events) event to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic and subscribe the topic to the HTTPS endpoint of the management application
- C. Modify the on-premises management application to call the AWS Health API to poll for status events of AWS services.
- D. Configure Amazon EventBridge (Amazon CloudWatch Events) to detect and react to changes for AWS Health events from the AWS Service Health Dashboard. Configure the EventBridge (CloudWatch Events) event to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic and subscribe the topic to an HTTPS endpoint for the management application with a topic filter corresponding to the services being used

Answer: A

Explanation:

ALB & NLB both supports IPs as targets. Questions is based on TCP traffic over VPN to on-premise. TCP is layer 4 and the , load balancer should be NLB. Then next questions does NLB supports loadbalancing traffic over VPN. And answer is YEs based on below URL.

<https://aws.amazon.com/about-aws/whats-new/2018/09/network-load-balancer-now-supports-aws-vpn/>

Target as IPs for NLB & ALB: <https://aws.amazon.com/elasticloadbalancing/faqs/?nc=sn&loc=5> <https://aws.amazon.com/elasticloadbalancing/application-load-balancer/>

NEW QUESTION 131

- (Exam Topic 2)

A company is building a software-as-a-service (SaaS) solution on AWS. The company has deployed an Amazon API Gateway REST API with AWS Lambda integration in multiple AWS Regions and in the same production account.

The company offers tiered pricing that gives customers the ability to pay for the capacity to make a certain number of API calls per second. The premium tier offers up to 3,000 calls per second, and customers are identified by a unique API key. Several premium tier customers in various Regions report that they receive error responses of 429 Too Many Requests from multiple API methods during peak usage hours. Logs indicate that the Lambda function is never invoked.

What could be the cause of the error messages for these customers?

- A. The Lambda function reached its concurrency limit.
- B. The Lambda function its Region limit for concurrency.
- C. The company reached its API Gateway account limit for calls per second.
- D. The company reached its API Gateway default per-method limit for calls per second.

Answer: C

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-request-throttling.html#apig-reques>

NEW QUESTION 136

- (Exam Topic 2)

A company manages multiple AWS accounts by using AWS Organizations. Under the root OU, the company has two OUs: Research and DataOps.

Because of regulatory requirements, all resources that the company deploys in the organization must reside in the ap-northeast-1 Region. Additionally, EC2 instances that the company deploys in the DataOps OU must use a predefined list of instance types

A solutions architect must implement a solution that applies these restrictions. The solution must maximize operational efficiency and must minimize ongoing maintenance

Which combination of steps will meet these requirements? (Select TWO)

- A. Create an IAM role in one account under the DataOps OU Use the ec2 Instance Type condition key in an inline policy on the role to restrict access to specific instance types.
- B. Create an IAM user in all accounts under the root OU Use the aws RequestedRegion condition key in an inline policy on each user to restrict access to all AWS Regions except ap-northeast-1.
- C. Create an SCP Use the aws:RequestedRegion condition key to restrict access to all AWS Regions except ap-northeast-1 Apply the SCP to the root OU.
- D. Create an SCP Use the ec2:InstanceType condition key to restrict access to all AWS Regions except ap-northeast-1. Apply the SCP to the root O
- E. the DataOps O
- F. and the Research OU.
- G. Create an SCP Use the ec2:InstanceType condition key to restrict access to specific instance types Apply the SCP to the DataOps OU.

Answer: CE

NEW QUESTION 140

- (Exam Topic 2)

A company plans to refactor a monolithic application into a modern application designed deployed or AWS. The CLCD pipeline needs to be upgraded to support the modern design for the application with the following requirements

- It should allow changes to be released several times every hour.
- * It should be able to roll back the changes as quickly as possible Which design will meet these requirements?

- A. Deploy a CI-CD pipeline that incorporates AMIs to contain the application and their configurations Deploy the application by replacing Amazon EC2 instances
- B. Specify AWS Elastic Beanstalk to sage in a secondary environment as the deployment target for the CI/CD pipeline of the applicatio
- C. To deploy swap the staging and production environment URLs.
- D. Use AWS Systems Manager to re-provision the infrastructure for each deployment Update the Amazon EC2 user data to pull the latest code art-fact from Amazon S3 and use Amazon Route 53 weighted routing to point to the new environment
- E. Roll out At application updates as pan of an Auto Scaling event using prebuilt AMI
- F. Use new versions of the AMIs to add instances, and phase out all instances that use the previous AMI version with the configured termination policy during a deployment event.

Answer: B

Explanation:

It is the fastest when it comes to rollback and deploying changes every hour

NEW QUESTION 142

- (Exam Topic 2)

A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company uses AWS Control Tower for governance and uses AWS Transit Gateway for VPC connectivity across accounts.

In an AWS application account, the company's application team has deployed a web application that uses AWS Lambda and Amazon RDS. The company's database administrators have a separate DBA account and use the account to centrally manage all the databases across the organization. The database administrators use an Amazon EC2 instance that is deployed in the DBA account to access an RDS database that is deployed in the application account.

The application team has stored the database credentials as secrets in AWS Secrets Manager in the application account. The application team is manually sharing the secrets with the database administrators. The secrets are encrypted by the default AWS managed key for Secrets Manager in the application account. A

solutions architect needs to implement a solution that gives the database administrators access to the database and eliminates the need to manually share the secrets.

Which solution will meet these requirements?

- A. Use AWS Resource Access Manager (AWS RAM) to share the secrets from the application account with the DBA account
- B. In the DBA account, create an IAM role that is named DBA-Admi
- C. Grant the role the required permissions to access the shared secret
- D. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
- E. In the application account, create an IAM role that is named DBA-Secre
- F. Grant the role the required permissions to access the secret
- G. In the DBA account, create an IAM role that is named DBA-Admi
- H. Grant the DBA-Admin role the required permissions to assume the DBA-Secret role in the application account
- I. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
- J. In the DBA account, create an IAM role that is named DBA-Admi
- K. Grant the role the required permissions to access the secrets and the default AWS managed key in the application account
- L. In the application account, attach resource-based policies to the key to allow access from the DBA account
- M. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
- N. In the DBA account, create an IAM role that is named DBA-Admi
- O. Grant the role the required permissions to access the secrets in the application account
- P. Attach an SCP to the application account to allow access to the secrets from the DBA account
- Q. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.

Answer: A

NEW QUESTION 147

- (Exam Topic 2)

A company has an organization in AWS Organizations that has a large number of AWS accounts. One of the AWS accounts is designated as a transit account and has a transit gateway that is shared with all of the other AWS accounts AWS Site-to-Site VPN connections are configured between all of the company's global offices and the transit account The company has AWS Config enabled on all of its accounts.

The company's networking team needs to centrally manage a list of internal IP address ranges that belong to the global offices Developers Will reference this list to gain access to applications securely.

Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Create a JSON file that is hosted in Amazon S3 and that lists all of the internal IP address ranges Configure an Amazon Simple Notification Service (Amazon SNS) topic in each of the accounts that can be involved when the JSON file is update
- B. Subscribe an AWS Lambda function to the SNS topic to update all relevant security group rules with the updated IP address ranges.
- C. Create a new AWS Config managed rule that contains all of the internal IP address ranges Use the rule to check the security groups in each of the accounts to ensure compliance with the list of IP address range
- D. Configure the rule to automatically remediate any noncompliant security group that is detected.
- E. In the transit account, create a VPC prefix list with all of the internal IP address range
- F. Use AWS Resource Access Manager to share the prefix list with all of the other account
- G. Use the shared prefix list to configure security group rules in the other accounts.
- H. In the transit account create a security group with all of the internal IP address range
- I. Configure the security groups in the other accounts to reference the transit account's security group by using a nested security group reference of `*<transit-account-id>/.sg-1a2b3c4d"`.

Answer: C

NEW QUESTION 152

- (Exam Topic 2)

A company's interactive web application uses an Amazon CloudFront distribution to serve images from an Amazon S3 bucket. Occasionally, third-party tools ingest corrupted images into the S3 bucket. This image corruption causes a poor user experience in the application later. The company has successfully implemented and tested Python logic to detect corrupt images.

A solutions architect must recommend a solution to integrate the detection logic with minimal latency between the ingestion and serving.

Which solution will meet these requirements?

- A. Use a Lambda@Edge function that is invoked by a viewer-response event.
- B. Use a Lambda@Edge function that is invoked by an origin-response event.
- C. Use an S3 event notification that invokes an AWS Lambda function.
- D. Use an S3 event notification that invokes an AWS Step Functions state machine.

Answer: A

NEW QUESTION 156

- (Exam Topic 2)

A company wants to send data from its on-premises systems to Amazon S3 buckets. The company created the S3 buckets in three different accounts. The company must send the data privately without the data traveling across the internet. The company has no existing dedicated connectivity to AWS

Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

- A. Establish a networking account in the AWS Cloud Create a private VPC in the networking account Set up an AWS Direct Connect connection with a private VIF between the on-premises environment and the private VPC
- B. Establish a networking account in the AWS Cloud Create a private VPC in the networking account Set up an AWS Direct Connect connection with a public VIF between the on-premises environment and the private VPC
- C. Create an Amazon S3 interface endpoint in the networking account
- D. Create an Amazon S3 gateway endpoint in the networking account
- E. Establish a networking account in the AWS Cloud
- F. Create a private VPC in the networking account Peer VPCs from the accounts that host the S3 buckets with the VPC in the network account

Answer: AD

NEW QUESTION 159

- (Exam Topic 2)

A video processing company has an application that downloads images from an Amazon S3 bucket, processes the images, stores a transformed image in a second S3 bucket, and updates metadata about the image in an Amazon DynamoDB table. The application is written in Node.js and runs by using an AWS Lambda function. The Lambda function is invoked when a new image is uploaded to Amazon S3.

The application ran without incident for a while. However, the size of the images has grown significantly. The Lambda function is now failing frequently with timeout errors. The function timeout is set to its maximum value. A solutions architect needs to refactor the application's architecture to prevent invocation failures. The company does not want to manage the underlying infrastructure.

Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

- A. Modify the application deployment by building a Docker image that contains the application code. Publish the image to Amazon Elastic Container Registry (Amazon ECR).
- B. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of AWS Fargate
- C. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.
- D. Create an AWS Step Functions state machine with a Parallel state to invoke the Lambda function. Increase the provisioned concurrency of the Lambda function.
- E. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of Amazon EC2. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.
- F. Modify the application to store images on Amazon Elastic File System (Amazon EFS) and to store metadata on an Amazon RDS DB instance
- G. Adjust the Lambda function to mount the EFS file share.

Answer: DE

NEW QUESTION 161

- (Exam Topic 2)

A company runs a proprietary stateless ETL application on an Amazon EC2 Linux instance. The application is a Linux binary, and the source code cannot be modified. The application is single-threaded, uses 2 GB of RAM, and is highly CPU intensive. The application is scheduled to run every 4 hours and runs for up to 20 minutes. A solutions architect wants to revise the architecture for the solution.

Which strategy should the solutions architect use?

- A. Use AWS Lambda to run the application
- B. Use Amazon CloudWatch Logs to invoke the Lambda function every 4 hours
- C. Use AWS Batch to run the application. Use an AWS Step Functions state machine to invoke the AWS Batch job every 4 hours
- D. Use AWS Fargate to run the application. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke the Fargate task every 4 hours
- E. Use Amazon EC2 Spot Instances to run the application. Use AWS CodeDeploy to deploy and run the application every 4 hours.

Answer: C

NEW QUESTION 165

- (Exam Topic 2)

A company has migrated an application from on-premises to AWS. The application frontend is a static website that runs on two Amazon EC2 instances behind an Application Load Balancer (ALB). The application backend is a Python application that runs on three EC2 instances behind another ALB. The EC2 instances are large, general purpose On-Demand Instances that were sized to meet the on-premises specifications for peak usage of the application.

The application averages hundreds of thousands of requests each month. However, the application is used mainly during lunchtime and receives minimal traffic during the rest of the day.

A solutions architect needs to optimize the infrastructure cost of the application without negatively affecting the application availability.

Which combination of steps will meet these requirements? (Choose two.)

- A. Change all the EC2 instances to compute optimized instances that have the same number of cores as the existing EC2 instances.
- B. Move the application frontend to a static website that is hosted on Amazon S3.
- C. Deploy the application frontend by using AWS Elastic Beanstalk
- D. Use the same instance type for the nodes.
- E. Change all the backend EC2 instances to Spot Instances.
- F. Deploy the backend Python application to general purpose burstable EC2 instances that have the same number of cores as the existing EC2 instances.

Answer: BE

NEW QUESTION 167

- (Exam Topic 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-intensive workloads such as simulations. The number of accounts changes frequently because of new initiatives, acquisitions, and divestitures.

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

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

- A. Store the data in a single Amazon S3 bucket. Create an IAM role for every combination of job type and business unit that allows for appropriate read/write access based on object prefixes in the S3 bucket. The roles should have trust policies that allow the business unit's AWS accounts to assume their roles. Use IAM in each business unit's AWS account to prevent them from assuming roles for a different job type. Users get credentials to access the data by using AssumeRole from their business unit's AWS account. Users can then use those credentials with an S3 client.
- B. Store the data in a single Amazon S3 bucket. 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.
- C. Determine the business unit with the AWS account accessing the bucket and the job type with a prefix in the IAM user's name. Users can access data by using IAM credentials from their business unit's AWS account with an S3 client.
- D. Store the data in a series of Amazon S3 buckets. Create an application running on 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. The application uses the business unit and job type information in the IdP to control what users can upload and download through the application. The users can access the data through the application's API.
- E. Store the data in a series of Amazon S3 buckets. 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 Users can get credentials by authenticating to the token vending machine's website or API and then use those credentials with an S3 client F. D

Answer: E

NEW QUESTION 169

- (Exam Topic 2)

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

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

Answer: A

NEW QUESTION 174

- (Exam Topic 2)

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

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

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

- A. Create two cache behaviors for static and dynamic content Remove the user-Agent and Host HTTP headers from the allow list headers section on both of the cache behaviors Remove the session cookie from the allow list cookies section and the Authorization HTTP header from the allow list headers section for cache behavior configured for static content
- B. Remove the user-Agent and Authorization HTTP headers from the allow list headers section of the cache behaviour
- C. Then update the cache behaviour to use resigned cookies for authorization
- D. Remove the Host HTTP header from the allow list headers section and remove the session cookie from the allow list cookies section for the default cache behaviour Enable automatic object compression and use Lambda@Edge viewer request events for user authorization
- E. Create two cache behaviours for static and dynamic content Remove the User-Agent HTTP header from the allow list headers section on both of the cache behavioursRemove the session cookie from the allow list cookies section and the Authorization HTTP header from the allow list headers section for cache behaviour configured for static content

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/understanding-the-cache-key.html> Removing the host header will result in failed flow between CloudFront and ALB, because they have same certificate.

NEW QUESTION 179

- (Exam Topic 2)

During an audit, a security team discovered that a development team was putting IAM user secret access keys in their code and then committing it to an AWS CodeCommit repository . The security team wants to automatically find and remediate instances of this security vulnerability Which solution will ensure that the credentials are appropriately secured automatically?

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

Answer: A

NEW QUESTION 182

- (Exam Topic 2)

A company runs its application in the eu-west-1 Region and has one account for each of its environments development, testing, and production All the environments are running 24 hours a day 7 days a week by using stateful Amazon EC2 instances and Amazon RDS for MySQL databases The databases are between 500 GB and 800 GB in size

The development team and testing team work on business days during business hours, but the production environment operates 24 hours a day. 7 days a week. The company wants to reduce costs AH resources are tagged with an environment tag with either development, testing, or production as the key.

What should a solutions architect do to reduce costs with the LEAST operational effort?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs once every day Configure the rule to invoke one AWS Lambda function that starts

or stops instances based on the tag day and time.

- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every business day in the evening
- C. Configure the rule to invoke an AWS Lambda function that stops instances based on the tag>Create a second EventBridge (CloudWatch Events) rule that runs every business day in the morning Configure the second rule to invoke another Lambda function that starts instances based on the tag
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every business day in the evening Configure the rule to invoke an AWS Lambda function that terminates instances based on the tag>Create a second EventBridge (CloudWatch Events) rule that runs every business day in the morning Configure the second rule to invoke another Lambda function that restores the instances from their last backup based on the tag.
- E. Create an Amazon EventBridge rule that runs every hour
- F. Configure the rule to invoke one AWS Lambda function that terminates or restores instances from their last backup based on the tag
- G. day, and time.

Answer: A

NEW QUESTION 186

- (Exam Topic 2)

A company is creating a sequel for a popular online game. A large number of users from all over the world will play the game within the first week after launch. Currently, the game consists of the following components deployed in a single AWS Region:

- Amazon S3 bucket that stores game assets
- Amazon DynamoDB table that stores player scores

A solutions architect needs to design a multi-Region solution that will reduce latency improve reliability, and require the least effort to implement

What should the solutions architect do to meet these requirements?

- A. Create an Amazon CloudFront distribution to serve assets from the S3 bucket Configure S3Cross-Region Replication>Create a new DynamoDB table in a new Region Use the new table as a replica target for DynamoDB global tables.
- B. Create an Amazon CloudFront distribution to serve assets from the S3 bucket
- C. Configure S3Same-Region Replication
- D. Create a new DynamoDB table in a new Region
- E. Configure asynchronous replication between the DynamoDB tables by using AWS Database Migration Service (AWS DMS) with change data capture (CDC)
- F. Create another S3 bucket in a new Region and configure S3 Cross-Region Replication between the buckets>Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets in each Region
- G. Configure DynamoDB global tables by enabling Amazon DynamoDB Streams, and add a replica table in a new Region.
- H. Create another S3 bucket in the same Region, and configure S3 Same-Region Replication between the buckets- Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets>Create a new DynamoDB table in a new Region Use the new table as a replica target for DynamoDB global tables.

Answer: C

NEW QUESTION 188

- (Exam Topic 2)

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

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAllActions",
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*"
    },
    {
      "Sid": "DenyCloudTrail",
      "Effect": "Deny",
      "Action": "cloudtrail:*",
      "Resource": "*"
    }
  ]
}
```

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

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

Answer: C

Explanation:

is incorrect - https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

"SCPs are similar to AWS Identity and Access Management (IAM) permission policies and use almost the same syntax. However, an SCP never grants permissions."

SCPs alone are not sufficient to granting permissions to the accounts in your organization. No permissions are granted by an SCP. An SCP defines a guardrail, or sets limits, on the actions that the account's administrator can delegate to the IAM users and roles in the affected accounts. The administrator must still attach identity-based or resource-based policies to IAM users or roles, or to the resources in your accounts to actually grant permissions. The effective permissions are the logical intersection between what is allowed by the SCP and what is allowed by the IAM and resource-based policies.

NEW QUESTION 191

- (Exam Topic 2)

A software company is using three AWS accounts for each of its 10 development teams. The company has developed an AWS CloudFormation standard VPC template that includes three NAT gateways. The template is added to each account for each team. The company is concerned that network costs will increase each time a new development team is added. A solutions architect must maintain the reliability of the company's solutions and minimize operational complexity. What should the solutions architect do to reduce the network costs while meeting these requirements?

- A. Create a single VPC with three NAT gateways in a shared services account. Configure each account VPC with a default route through a transit gateway to the NAT gateway in the shared services account VPC. Remove all NAT gateways from the standard VPC template.
- B. Create a single VPC with three NAT gateways in a shared services account. Configure each account VPC with a default route through a VPC peering connection to the NAT gateway in the shared services account VPC. Remove all NAT gateways from the standard VPC template.
- C. Remove two NAT gateways from the standard VPC template. Rely on the NAT gateway SLA to cover reliability for the remaining NAT gateway.
- D. Create a single VPC with three NAT gateways in a shared services account. Configure a Site-to-Site VPN connection from each account to the shared services account. Remove all NAT gateways from the standard VPC template.

Answer: A

NEW QUESTION 194

- (Exam Topic 2)

A company has a platform that contains an Amazon S3 bucket for user content. The S3 bucket has thousands of terabytes of objects, all in the S3 Standard storage class. The company has an RTO of 6 hours. The company must replicate the data from its primary AWS Region to a replication S3 bucket in another Region.

The user content S3 bucket contains user-uploaded files such as videos and photos. The user content S3 bucket has an unpredictable access pattern. The number of users is increasing quickly, and the company wants to create an S3 Lifecycle policy to reduce storage costs. Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

- A. Move the objects in the user content S3 bucket to S3 Intelligent-Tiering immediately.
- B. Move the objects in the user content S3 bucket to S3 Intelligent-Tiering after 30 days.
- C. Move the objects in the replication S3 bucket to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days and to S3 Glacier after 90 days.
- D. Move the objects in the replication S3 bucket to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days and to S3 Glacier Deep Archive after 90 days.
- E. Move the objects in the replication S3 bucket to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days and to S3 Glacier Deep Archive after 180 days.

Answer: AD

NEW QUESTION 198

- (Exam Topic 2)

A solutions architect uses AWS Organizations to manage several AWS accounts for a company. The full Organizations feature set is activated for the organization. All production AWS accounts exist under an OU that is named "production." Systems operators have full administrative privileges within these accounts by using IAM roles.

The company wants to ensure that security groups in all production accounts do not allow inbound traffic for TCP port 22. All noncompliant security groups must be remediated immediately, and no new rules that allow port 22 can be created.

Which solution will meet these requirements?

- A. Write an SCP that denies the CreateSecurityGroup action with a condition of `ec2:ingress rule with value 22`. Apply the SCP to the 'production' OU.
- B. Configure an AWS CloudTrail trail for all accounts. Send CloudTrail logs to an Amazon S3 bucket in the Organizations management account.
- C. Configure an AWS Lambda function on the management account with permissions to assume a role in all production accounts to describe and modify security groups.
- D. Configure Amazon S3 to invoke the Lambda function on every PutObject event on the S3 bucket. Configure the Lambda function to analyze each CloudTrail event for noncompliant security group actions and to automatically remediate any issues.
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) event bus in the Organizations management account.
- F. Create an AWS CloudFormation template to deploy configurations that send CreateSecurityGroup events to the event bus from all production accounts. Configure an AWS Lambda function in the management account with permissions to assume a role in all production accounts to describe and modify security groups.
- G. Configure the event bus to invoke the Lambda function. Configure the Lambda function to analyze each event for noncompliant security group actions and to automatically remediate any issues.
- H. Create an AWS CloudFormation template to turn on AWS Config. Activate the INCOMING_SSH_DISABLED AWS Config managed rule. Deploy an AWS Lambda function that will run based on AWS Config findings and will remediate noncompliant resources. Deploy the CloudFormation template by using a StackSet that is assigned to the "production" OU.
- I. Apply an SCP to the OU to deny modification of the resources that the CloudFormation template provisions.

Answer: D

NEW QUESTION 202

- (Exam Topic 2)

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

NEW QUESTION 207

- (Exam Topic 2)

A company manages hundreds of AWS accounts centrally in an organization. In AWS Organizations, the company recently started to allow product teams to

create and manage their own S3 access points in their accounts. The S3 access points can be accessed only within VPCs, not on the internet. What is the MOST operationally efficient way to enforce this requirement?

- A. Set the S3 access point resource policy to deny the s3CreateAccessPoint action unless the s3 AccessPointNetworkOrigin condition key evaluates to VPC.
- B. Create an SCP at the root level in the organization to deny the s3: Create Access Point action unless the s3:AccessPointNetworkOrigin condition key evaluates to VPC.
- C. Use AWS Cloud Formation StackSets to create a new IAM policy in each AWS account that allows the s3:CreateAccessPoint action only if the s3:AccessPointNetworkOrigin condition key evaluates to VPC.
- D. Set the S3 bucket policy to deny the s3:CreateAccessPoint action unless the s3: AccessPointNetworkOrigin condition key evaluates to VPC.

Answer: A

NEW QUESTION 211

- (Exam Topic 2)

A company has an asynchronous HTTP application that is hosted as an AWS Lambda function. A public Amazon API Gateway endpoint invokes the Lambda function. The Lambda function and the API Gateway endpoint reside in the us-east-1 Region. A solutions architect needs to redesign the application to support failover to another AWS Region.

Which solution will meet these requirements?

- A. Create an API Gateway endpoint in the us-west-2 Region to direct traffic to the Lambda function in us-east-1. Configure Amazon Route 53 to use a failover routing policy to route traffic for the two API Gateway endpoints.
- B. Create an Amazon Simple Queue Service (Amazon SQS) queue.
- C. Configure API Gateway to direct traffic to the SQS queue instead of to the Lambda function.
- D. Configure the Lambda function to pull messages from the queue for processing.
- E. Deploy the Lambda function to the us-west-2 Region.
- F. Create an API Gateway endpoint in us-west-2 to direct traffic to the Lambda function in us-west-2. Configure AWS Global Accelerator and an Application Load Balancer to manage traffic across the two API Gateway endpoints.
- G. Deploy the Lambda function and an API Gateway endpoint to the us-west-2 Region.
- H. Configure Amazon Route 53 to use a failover routing policy to route traffic for the two API Gateway endpoints.

Answer: B

NEW QUESTION 212

- (Exam Topic 2)

A company is planning to migrate 1,000 on-premises servers to AWS. The servers run on several VMware clusters in the company's data center. As part of the migration plan, the company wants to gather server metrics such as CPU details, RAM usage, operating system information, and running processes. The company then wants to query and analyze the data.

Which solution will meet these requirements?

- A. Deploy and configure the AWS Agentless Discovery Connector virtual appliance on the on-premises host.
- B. Configure Data Exploration in AWS Migration Hub.
- C. Use AWS Glue to perform an ETL job against the data.
- D. Query the data by using Amazon S3 Select.
- E. Export only the VM performance information from the on-premises host.
- F. Directly import the required data into AWS Migration Hub.
- G. Update any missing information in Migration Hub.
- H. Query the data by using Amazon QuickSight.
- I. Create a script to automatically gather the server information from the on-premises host.
- J. Use the AWS CLI to run the put-resource-attributes command to store the detailed server data in AWS Migration Hub.
- K. Query the data directly in the Migration Hub console.
- L. Deploy the AWS Application Discovery Agent to each on-premises server.
- M. Configure Data Exploration in AWS Migration Hub.
- N. Use Amazon Athena to run predefined queries against the data in Amazon S3.

Answer: C

NEW QUESTION 215

.....

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

AWS-Certified-Solutions-Architect-Professional Practice Exam Features:

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

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Certified-Solutions-Architect-Professional Practice Test Here](#)