



Vendor: Amazon

Exam Code: SAP-C02

Exam Name: AWS Certified Solutions Architect -
Professional (SAP-C02)

Version: DEMO

QUESTION 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 account. Outbound internet traffic will be controlled by firewall appliances.
- B. A centralized shared VPC with a subnet for each account. Outbound internet traffic will be controlled through a fleet of proxy servers.
- C. 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.
- D. A shared transit gateway to which each VPC will be attached. 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/centralized-egress-to-internet.html>

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.

QUESTION 2

A solutions architect is responsible for redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application runs on a single high-memory Amazon EC2 instance. It accepts http requests from upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average, most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel. Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backlog. In addition, the current system has issues with availability and data loss if the single application node fails. Clients that access this service cannot be modified. They expect to receive a response to each http request they send within 10 seconds before they will time out and retry the request. Which approach would improve the availability and durability of the system while decreasing the processing latency and minimizing costs?

- A. Create an Amazon API Gateway REST API that uses Lambda proxy integration to pass requests to an AWS Lambda function. Migrate the core processing code to a Lambda function and write a wrapper class that provides a handler method that converts the proxy events to the internal application data model and invokes the processing module.
- B. Create an Amazon API Gateway REST API that uses a service proxy to put items in an Amazon SQS queue. Extract the core processing code from the existing application and update it to pull items from Amazon SQS instead of an in-memory queue. Deploy the new processing application to smaller EC2 instances within an Auto Scaling group that scales dynamically based on the approximate number of messages in the Amazon SQS queue.
- C. Modify the application to use Amazon DynamoDB instead of Amazon RDS. Configure Auto Scaling for the DynamoDB table. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilization. Back the in-memory queue with a memory-mapped file to

an instance store volume and periodically write that file to Amazon S3.

- D. Update the application to use a Redis task queue instead of the in-memory queue. Build a Docker container image for the application. Create an Amazon ECS task definition that includes the application container and a separate container to host Redis. Deploy the new task definition as an ECS service using AWS Fargate, and enable Auto Scaling.

Answer: B

Explanation:

The obvious challenges here are long workloads, scalability based on queue load, and reliability. Almost always the defacto answer to queue related workload is SQS. Since the workloads are very long (90 minutes) Lambdas cannot be used (15 mins max timeout). So, autoscaled smaller EC2 nodes that wait on external services to complete the task makes more sense. If the task fails, the message is returned to the queue and retried.

QUESTION 3

A company runs an application on AWS. An AWS Lambda function uses credentials to authenticate to an Amazon RDS for MySQL DB instance. A security risk assessment identified that these credentials are not frequently rotated. Also, encryption at rest is not enabled for the DB instance. The security team requires that both of these issues be resolved.

Which strategy should a solutions architect recommend to remediate these security risks?

- A. Configure the Lambda function to store and retrieve the database credentials in AWS Secrets Manager and enable rotation of the credentials.
Take a snapshot of the DB instance and encrypt a copy of that snapshot.
Replace the DB instance with a new DB instance that is based on the encrypted snapshot.
- B. Enable IAM DB authentication on the DB instance.
Grant the Lambda execution role access to the DB instance.
Modify the DB instance and enable encryption.
- C. Enable IAM DB authentication on the DB instance.
Grant the Lambda execution role access to the DB instance.
Create an encrypted read replica of the DB instance.
Promote the encrypted read replica to be the new primary node.
- D. Configure the Lambda function to store and retrieve the database credentials as encrypted AWS Systems Manager Parameter Store parameters.
Create another Lambda function to automatically rotate the credentials.
Create an encrypted read replica of the DB instance.
Promote the encrypted read replica to be the new primary node.

Answer: A

Explanation:

Parameter store can store DB credentials as secure string but CANNOT rotate secrets, hence, go with A + Cannot enable encryption on existing MySQL RDS instance, must create a new encrypted one from unencrypted snapshot.

<https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html>

QUESTION 4

A fitness tracking company serves users around the world, with its primary markets in North America and Asia. The company needs to design an infrastructure for its read-heavy user authorization application with the following requirements:

- Be resilient to problems with the application in any Region.
- Write to a database in a single Region.
- Read from multiple Regions.

- Support resiliency across application tiers in each Region.
- Support the relational database semantics reflected in the application.

Which combination of steps should a solutions architect take? (Choose two.)

- A. Use an Amazon Route 53 geoproximity routing policy combined with a multivalue answer routing policy.
- B. Deploy web, application, and MySQL database servers to Amazon EC2 instances in each Region.
Set up the application so that reads and writes are local to the Region. Create snapshots of the web, application, and database servers and store the snapshots in an Amazon S3 bucket in both Regions.
Set up cross-Region replication for the database layer.
- C. Use an Amazon Route 53 geolocation routing policy combined with a failover routing policy.
- D. Set up web, application, and Amazon RDS for MySQL instances in each Region. Set up the application so that reads are local and writes are partitioned based on the user. Set up a Multi-AZ failover for the web, application, and database servers. Set up cross-Region replication for the database layer.
- E. Set up active-active web and application servers in each Region. Deploy an Amazon Aurora global database with clusters in each Region. Set up the application to use the in-Region Aurora database endpoints. Create snapshots of the web and application servers and store them in an Amazon S3 bucket in both Regions.

Answer: C

Explanation:

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

Geoproximity routing policy is good to control the user traffic to specific regions. However, a multivalue answer routing policy may cause the users to be randomly sent to other healthy regions that may be far away from the user's location. You can use geolocation routing policy to direct the North American users to your servers on the North America region and configure failover routing to the Asia region in case the North America region fails. You can configure the same for the Asian users pointed to the Asia region servers and have the North America region as its backup.

QUESTION 5

A company has a data lake in Amazon S3 that needs to be accessed by hundreds of applications across many AWS accounts. The company's information security policy states that the S3 bucket must not be accessed over the public internet and that each application should have the minimum permissions necessary to function.

To meet these requirements, a solutions architect plans to use an S3 access point that is restricted to specific VPCs for each application.

Which combination of steps should the solutions architect take to implement this solution? (Choose two.)

- A. Create an S3 access point for each application in the AWS account that owns the S3 bucket. Configure each access point to be accessible only from the application's VPC. Update the bucket policy to require access from an access point.
- B. Create an interface endpoint for Amazon S3 in each application's VPC. Configure the endpoint policy to allow access to an S3 access point. Create a VPC gateway attachment for the S3 endpoint.
- C. Create a gateway endpoint for Amazon S3 in each application's VPC. Configure the endpoint policy to allow access to an S3 access point. Specify the route table that is used to access the access point.

- D. Create an S3 access point for each application in each AWS account and attach the access points to the S3 bucket.
Configure each access point to be accessible only from the application's VPC.
Update the bucket policy to require access from an access point.
- E. Create a gateway endpoint for Amazon S3 in the data lake's VPC.
Attach an endpoint policy to allow access to the S3 bucket.
Specify the route table that is used to access the bucket.

Answer: AC

Explanation:

App --> S3 Access Point --> S3 Gateway Endpoint --> S3 Bucket

<https://aws.amazon.com/blogs/storage/managing-amazon-s3-access-with-vpc-endpoints-and-s3-access-points/>

QUESTION 6

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-provisioned-iops-for-your-rds-database/>

QUESTION 7

A company is creating a REST API to share information with six of its partners based in the United States.

The company has created an Amazon API Gateway Regional endpoint.

Each of the six partners will access the API once per day to post daily sales figures.

After initial deployment the company observes 1,000 requests per second originating from 500 different IP addresses around the world.

The company believes this traffic is originating from a botnet and wants to secure its API while minimizing cost.

Which approach should the company take to secure its API?

- A. Create an Amazon CloudFront distribution with the API as the origin.
Create an AWS WAF web ACL with a rule to block clients that submit more than five requests per

day.

Associate the web ACL with the CloudFront distribution.

Configure CloudFront with an origin access identity (OAI) and associate it with the distribution.

Configure API Gateway to ensure only the OAI can execute the POST method.

- B. Create an Amazon CloudFront distribution with the API as the origin.

Create an AWS WAF web ACL with a rule to block clients that submit more than five requests per day.

Associate the web ACL with the CloudFront distribution.

Add a custom header to the CloudFront distribution populated with an API key.

Configure the API to require an API key on the POST method.

- C. Create an AWS WAF web ACL with a rule to allow access to the IP addresses used by the six partners.

Associate the web ACL with the API.

Create a resource policy with a request limit and associate it with the API.

Configure the API to require an API key on the POST method.

- D. Create an AWS WAF web ACL with a rule to allow access to the IP addresses used by the six partners.

Associate the web ACL with the API.

Create a usage plan with a request limit and associate it with the API.

Create an API key and add it to the usage plan.

Answer: D

Explanation:

A usage plan specifies who can access one or more deployed API stages and methods - and also how much and how fast they can access them. The plan uses API keys to identify API clients and meters access to the associated API stages for each key. It also lets you configure throttling limits and quota limits that are enforced on individual client API keys.

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html>

QUESTION 8

A company that tracks medical devices in hospitals wants to migrate its existing storage solution to the AWS Cloud. The company equips all of its devices with sensors that collect location and usage information. This sensor data is sent in unpredictable patterns with large spikes. The data is stored in a MySQL database running on premises at each hospital. The company wants the cloud storage solution to scale with usage.

The company's analytics team uses the sensor data to calculate usage by device type and hospital. The team needs to keep analysis tools running locally while fetching data from the cloud. The team also needs to use existing Java application and SQL queries with as few changes as possible. How should a solutions architect meet these requirements while ensuring the sensor data is secure?

- A. Store the data in an Amazon Aurora Serverless database. Serve the data through a Network Load Balancer (NLB). Authenticate users using the NLB with credentials stored in AWS Secrets Manager.
- B. Store the data in an Amazon S3 bucket. Serve the data through Amazon QuickSight using an IAM user authorized with AWS Identity and Access Management (IAM) with the S3 bucket as the data source.
- C. Store the data in an Amazon Aurora Serverless database. Serve the data through the Aurora Data API using an IAM user authorized with AWS Identity and Access Management (IAM) and the AWS Secrets Manager ARN.
- D. Store the data in an Amazon S3 bucket. Serve the data through Amazon Athena using AWS PrivateLink to secure the data in transit.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/new-data-api-for-amazon-aurora-serverless/>
<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/data-api.html>
<https://aws.amazon.com/blogs/aws/aws-privatelink-for-amazon-s3-now-available/>
<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/data-api.html#data-api.access>

The data is currently stored in a MySQL database running on-prem. Storing MySQL data in S3 doesn't sound good so B & D are out. Aurora Data API "enables the SQL http endpoint, a connectionless Web Service API for running SQL queries against this database. When the SQL http endpoint is enabled, you can also query your database from inside the RDS console (these features are free to use)."

QUESTION 9

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 Balancer. Create additional read replicas for the DB instance. Create Amazon Kinesis data streams and configure the application services to use the data streams. Store and serve static content directly from Amazon S3.
- B. Create an EC2 Auto Scaling group behind an Application Load Balancer. Deploy the DB instance in Multi-AZ mode and enable storage auto scaling. Create Amazon Kinesis data streams and configure the application services to use the data streams. Store and serve static content directly from Amazon S3.
- C. Deploy the application on a Kubernetes cluster created on the EC2 instances behind an Application Load Balancer. Deploy the DB instance in Multi-AZ mode and enable storage auto scaling. 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.
- D. 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.

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.

QUESTION 10

A medical company is running a REST API on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group behind an Application Load Balancer (ALB). The ALB runs in three public subnets, and the EC2 instances run in three private subnets. The company has deployed an Amazon CloudFront distribution that has the ALB as the only origin.

Which solution should a solutions architect recommend to enhance the origin security?

- A. Store a random string in AWS Secrets Manager. Create an AWS Lambda function for automatic secret rotation. Configure CloudFront to inject the random string as a custom http header for the origin request. Create an AWS WAF web ACL rule with a string match rule for the custom header. Associate the web ACL with the ALB.
- B. Create an AWS WAF web ACL rule with an IP match condition of the CloudFront service IP address ranges. Associate the web ACL with the ALB. Move the ALB into the three private subnets.
- C. Store a random string in AWS Systems Manager Parameter Store. Configure Parameter Store automatic rotation for the string. Configure CloudFront to inject the random string as a custom http header for the origin request. Inspect the value of the custom http header, and block access in the ALB.
- D. Configure AWS Shield Advanced. Create a security group policy to allow connections from CloudFront service IP address ranges. Add the policy to AWS Shield Advanced, and attach the policy to the ALB.

Answer: D

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html>
It shows For Amazon EC2 Auto Scaling, there are two primary process types: Launch and Terminate. The Launch process adds a new Amazon EC2 instance to an Auto Scaling group, increasing its capacity. The Terminate process removes an Amazon EC2 instance from the group, decreasing its capacity. HealthCheck process for EC2 autoscaling is not a primary process! It is a process along with the following AddToLoadBalancer AlarmNotification AZRebalance HealthCheck InstanceRefresh ReplaceUnhealthy ScheduledActions From the requirements, Some EC2 instances are now being marked as unhealthy and are being terminated. Application is running at reduced capacity not because instances are marked unhealthy but because they are being terminated.
<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html#choosing-suspend-resume>

QUESTION 11

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. Attach the shared VPC and the authorized business unit VPCs to the transit gateway. Create a single transit gateway route table and associate it with all of the attached VPCs. Allow automatic propagation of routes from the attachments into the route table. Configure VPC routing tables to send traffic to the transit gateway.
- B. Create a VPC endpoint service using the centralized application NLB and enable (the option to require endpoint acceptance). Create a VPC endpoint in each of the business unit VPCs using the service name of the endpoint service. Accept authorized endpoint requests from the endpoint service console.
- C. Create a VPC peering connection from each business unit VPC to the shared VPC. Accept the VPC peering connections from the shared VPC console. Configure VPC routing tables to send

traffic to the VPC peering connection.

- D. Configure a virtual private gateway for the shared VPC and create customer gateways for each of the authorized business unit VPCs. Establish a Site-to-Site VPN connection from the business unit VPCs to the shared VPC. 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-preservation>

QUESTION 12

A solutions architect is designing a publicly accessible web application that is on an Amazon CloudFront distribution with an Amazon S3 website endpoint as the origin. When the solution is deployed, the website returns an Error 403: Access Denied message.

Which steps should the solutions architect take to correct the issue? (Choose two.)

- A. Remove the S3 block public access option from the S3 bucket.
- B. Remove the requester pays option from the S3 bucket.
- C. Remove the origin access identity (OAI) from the CloudFront distribution.
- D. Change the storage class from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA).
- E. Disable S3 object versioning.

Answer: AB

Explanation:

See using S3 to host a static website with Cloudfront:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-serve-static-website/>

- Using a REST API endpoint as the origin, with access restricted by an origin access identity (OAI)
- Using a website endpoint as the origin, with anonymous (public) access allowed
- Using a website endpoint as the origin, with access restricted by a Referer header

QUESTION 13

A start up company hosts a fleet of Amazon EC2 instances in private subnets using the latest Amazon Linux 2 AMI. The company's engineers rely heavily on SSH access to the instances for troubleshooting.

The company's existing architecture includes the following:

- A VPC with private and public subnets, and a NAT gateway
- Site-to-Site VPN for connectivity with the on-premises environment
- EC2 security groups with direct SSH access from the on-premises environment

The company needs to increase security controls around SSH access and provide auditing of commands executed by the engineers.

Which strategy should a solutions architect use?

- A. Install and configure EC2 Instance Connect on the fleet of EC2 instances. Remove all security group rules attached to EC2 instances that allow inbound TCP on port 22. Advise the engineers

- to remotely access the instances by using the EC2 Instance Connect CLI.
- B. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's devices. Install the Amazon CloudWatch agent on all EC2 instances and send operating system audit logs to CloudWatch Logs.
 - C. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's devices. Enable AWS Config for EC2 security group resource changes. Enable AWS Firewall Manager and apply a security group policy that automatically remediates changes to rules.
 - D. Create an IAM role with the AmazonSSMManagedInstanceCore managed policy attached. Attach the IAM role to all the EC2 instances. Remove all security group rules attached to the EC2 instances that allow inbound TCP on port 22. Have the engineers install the AWS Systems Manager Session Manager plugin for their devices and remotely access the instances by using the start-session API call from Systems Manager.

Answer: D

Explanation:

Allows client machines to be able to connect to Session Manager using the AWS CLI instead of going through the AWS EC2 or AWS Server Manager console.

<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-working-with-install-plugin.html>

QUESTION 14

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

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

Answer: D

Explanation:

This feature also allows you to connect to any of the participating VPCs from any Direct Connect location, further reducing your costs for making using AWS services on a cross-region basis.

<https://aws.amazon.com/blogs/aws/new-aws-direct-connect-gateway-inter-region-vpc-access/>
<https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect-aws-transit-gateway.html>

QUESTION 15

A company needs to architect a hybrid DNS solution.

This solution will use an Amazon Route 53 private hosted zone for the domain

cloud.example.com for the resource stored within VPCs.
The company has the following DNS resolution requirements:

- On-premises systems should be able to resolve and connect to cloud.example.com.
- All VPCs should be able to resolve cloud.example.com.

There is already an AWS Direct Connect connection between the on-premises corporate network and AWS Transit Gateway.

Which architecture should the company use to meet these requirements with the HIGHEST performance?

- A. Associate the private hosted zone to all the VPCs.
Create a Route 53 inbound resolver in the shared services VPC.
Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.
- B. Associate the private hosted zone to all the VPCs.
Deploy an Amazon EC2 conditional forwarder in the shared services VPC.
Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the conditional forwarder.
- C. Associate the private hosted zone to the shared services VPC.
Create a Route 53 outbound resolver in the shared services VPC.
Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the outbound resolver.
- D. Associate the private hosted zone to the shared services VPC.
Create a Route 53 inbound resolver in the shared services VPC.
Attach the shared services VPC to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.

Answer: A

Explanation:

When a Route 53 private hosted zone needs to be resolved in multiple VPCs and AWS accounts as described earlier, the most reliable pattern is to share the private hosted zone between accounts and associate it to each VPC that needs it.

<https://aws.amazon.com/blogs/networking-and-content-delivery/centralized-dns-management-of-hybrid-cloud-with-amazon-route-53-and-aws-transit-gateway/>

QUESTION 16

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

Answer: A

Explanation:

You can use AWS Systems Manager Configuration Compliance to scan your fleet of managed instances for patch compliance.

Reference:

<https://aws.amazon.com/blogs/mt/how-moodys-uses-aws-systems-manager-to-patch-servers-across-multiple-cloud-providers/>

QUESTION 17

An online e-commerce business is running a workload on AWS. The application architecture includes a web tier, an application tier for business logic, and a database tier for user and transactional data management. The database server has a 100 GB memory requirement. The business requires cost-efficient disaster recovery for the application with an RTO of 5 minutes and an RPO of 1 hour. The business also has a regulatory requirement for out-of-region disaster recovery with a minimum distance between the primary and alternate sites of 250 miles. Which of the following options can the solutions architect design to create a comprehensive solution for this customer that meets the disaster recovery requirements?

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

Answer: C

Explanation:

As RTO is in minutes

<https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/plan-for-disaster-recovery-dr.html>

Warm standby (RPO in seconds, RTO in minutes): Maintain a scaled-down version of a fully functional environment always running in the DR Region. Business-critical systems are fully duplicated and are always on, but with a scaled down fleet. When the time comes for recovery, the system is scaled up quickly to handle the production load.

QUESTION 18

A company maintains a restaurant review website. The website is a single-page application where files are stored in Amazon S3 and delivered using Amazon CloudFront. The company receives several fake postings every day that are manually removed.

The security team has identified that most of the fake posts are from bots with IP addresses that have a bad reputation within the same global region. The team needs to create a solution to help restrict the bots from accessing the website.

Which strategy should a solutions architect use?

- A. Use AWS Firewall Manager to control the CloudFront distribution security settings. Create a geographical block rule and associate it with Firewall Manager.
- B. Associate an AWS WAF web ACL with the CloudFront distribution. Select the managed Amazon IP reputation rule group for the web ACL with a deny action.
- C. Use AWS Firewall Manager to control the CloudFront distribution security settings. Select the managed Amazon IP reputation rule group and associate it with Firewall Manager with a deny action.
- D. Associate an AWS WAF web ACL with the CloudFront distribution. Create a rule group for the web ACL with a geographical match statement with a deny action.

Answer: B

Explanation:

IP reputation rule groups allow you to block requests based on their source. Choose one or more of these rule groups if you want to reduce your exposure to BOTS!!!! traffic or exploitation attempts. The Amazon IP reputation list rule group contains rules that are based on Amazon internal threat intelligence. This is useful if you would like to block IP addresses typically associated with bots or other threats. Inspects for a list of IP addresses that have been identified as bots by Amazon threat intelligence.

QUESTION 19

A company has an application that generates reports and stores them in an Amazon S3 bucket. When a user accesses their report, the application generates a signed URL to allow the user to download the report. The company's security team has discovered that the files are public and that anyone can download them without authentication. The company has suspended the generation of new reports until the problem is resolved.

Which set of actions will immediately remediate the security issue without impacting the application's normal workflow?

- A. Create an AWS Lambda function that applies a deny all policy for users who are not authenticated.
Create a scheduled event to invoke the Lambda function.
- B. Review the AWS Trusted Advisor bucket permissions check and implement the recommended actions.
- C. Run a script that puts a private ACL on all of the objects in the bucket.
- D. Use the Block Public Access feature in Amazon S3 to set the IgnorePublicAcls option to TRUE on the bucket.

Answer: D

Explanation:

The S3 bucket is allowing public access and this must be immediately disabled. Setting the IgnorePublicAcls option to TRUE causes Amazon S3 to ignore all public ACLs on a bucket and any objects that it contains. The other settings you can configure with the Block Public Access Feature are:

BlockPublicAcls - PUT bucket ACL and PUT objects requests are blocked if granting public access.

BlockPublicPolicy - Rejects requests to PUT a bucket policy if granting public access.

RestrictPublicBuckets - Restricts access to principles in the bucket owners' AWS account.

<https://aws.amazon.com/s3/features/block-public-access/>

QUESTION 20

A company has multiple AWS accounts as part of an organization created with AWS

Organizations. Each account has a VPC in the us-east-2 Region and is used for either production or development workloads. Amazon EC2 instances across production accounts need to communicate with each other, and EC2 instances across development accounts need to communicate with each other, but production and development instances should not be able to communicate with each other. To facilitate connectivity, the company created a common network account. The company used AWS Transit Gateway to create a transit gateway in the us-east-2 Region in the network account and shared the transit gateway with the entire organization by using AWS Resource Access Manager. Network administrators then attached VPCs in each account to the transit gateway, after which the EC2 instances were able to communicate across accounts. However, production and development accounts were also able to communicate with one another. Which set of steps should a solutions architect take to ensure production traffic and development traffic are completely isolated?

- A. Modify the security groups assigned to development EC2 instances to block traffic from production EC2 instances. Modify the security groups assigned to production EC2 instances to block traffic from development EC2 instances.
- B. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attached. Using the Network Manager feature of AWS Transit Gateway, create policies that restrict traffic between VPCs based on the value of this tag.
- C. Create separate route tables for production and development traffic. Delete each account's association and route propagation to the default AWS Transit Gateway route table. Attach development VPCs to the development AWS Transit Gateway route table and production VPCs to the production route table, and enable automatic route propagation on each attachment.
- D. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attached. Modify the AWS Transit Gateway routing table to route production tagged attachments to one another and development tagged attachments to one another.

Answer: C

Explanation:

The production and development route tables and VPC should be disassociated and configured separately using AWS transit gateway then route propagation enabled for the respective VPC attachments.

<https://aws.amazon.com/cn/blogs/architecture/field-notes-working-with-route-tables-in-aws-transit-gateway/>

QUESTION 21

A startup company recently migrated a large ecommerce website to AWS. The website has experienced a 70% increase in sales.

Software engineers are using a private GitHub repository to manage code. The development team is using Jenkins for builds and unit testing.

The engineers need to receive notifications for bad builds and zero downtime during deployments.

The engineers also need to ensure any changes to production are seamless for users and can be rolled back in the event of a major issue.

The software engineers have decided to use AWS CodePipeline to manage their build and deployment process.

Which solution will meet these requirements?

- A. Use GitHub websockets to trigger the CodePipeline pipeline.
Use the Jenkins plugin for AWS CodeBuild to conduct unit testing. Send alerts to an Amazon SNS topic for any bad builds. Deploy in an in-place all-at-once deployment configuration using AWS CodeDeploy.
- B. Use GitHub webhooks to trigger the CodePipeline pipeline.

- Use the Jenkins plugin for AWS CodeBuild to conduct unit testing. Deploy in a blue/green deployment using AWS CodeDeploy.
- C. Use GitHub websockets to trigger the CodePipeline pipeline.
Use AWS X-Ray for unit testing and static code analysis.
Send alerts to an Amazon SNS topic for any bad builds.
Deploy in a blue/green deployment using AWS CodeDeploy
- D. Use GitHub websockets to trigger the CodePipeline pipeline.
Use AWS X-Ray for unit testing and static code analysis.
Deploy in an in-place, all-at-once deployment configuration using AWS CodeDeploy

Answer: B

Explanation:

Use GitHub webhooks to trigger the CodePipeline pipeline. Use the Jenkins plugin for AWS CodeBuild to conduct unit testing. Send alerts to an Amazon SNS topic for any bad builds. Deploy in a blue/green deployment using AWS CodeDeploy.

QUESTION 22

A company wants to change its internal cloud billing strategy for each of its business units. Currently, the cloud governance team shares reports for overall cloud spending with the head of each business unit. The company uses AWS Organizations to manage the separate AWS accounts for each business unit. The existing tagging standard in Organizations includes the application, environment, and owner. The cloud governance team wants a centralized solution so each business unit receives monthly reports on its cloud spending. The solution should also send notifications for any cloud spending that exceeds a set threshold.

Which solution is the MOST cost-effective way to meet these requirements?

- A. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in each account to create monthly reports for each business unit.
- B. Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.
- C. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use the AWS Billing and Cost Management dashboard in each account to create monthly reports for each business unit.
- D. Enable AWS Cost and Usage Reports in the organization's master account and configure reports grouped by application, environment, and owner. Create an AWS Lambda function that processes AWS Cost and Usage Reports, sends budget alerts, and sends monthly reports to each business unit's email list.

Answer: B

Explanation:

Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.lead2pass.com/all-products.html>



10% Discount Coupon Code: ASTR14