PassleaderVCE

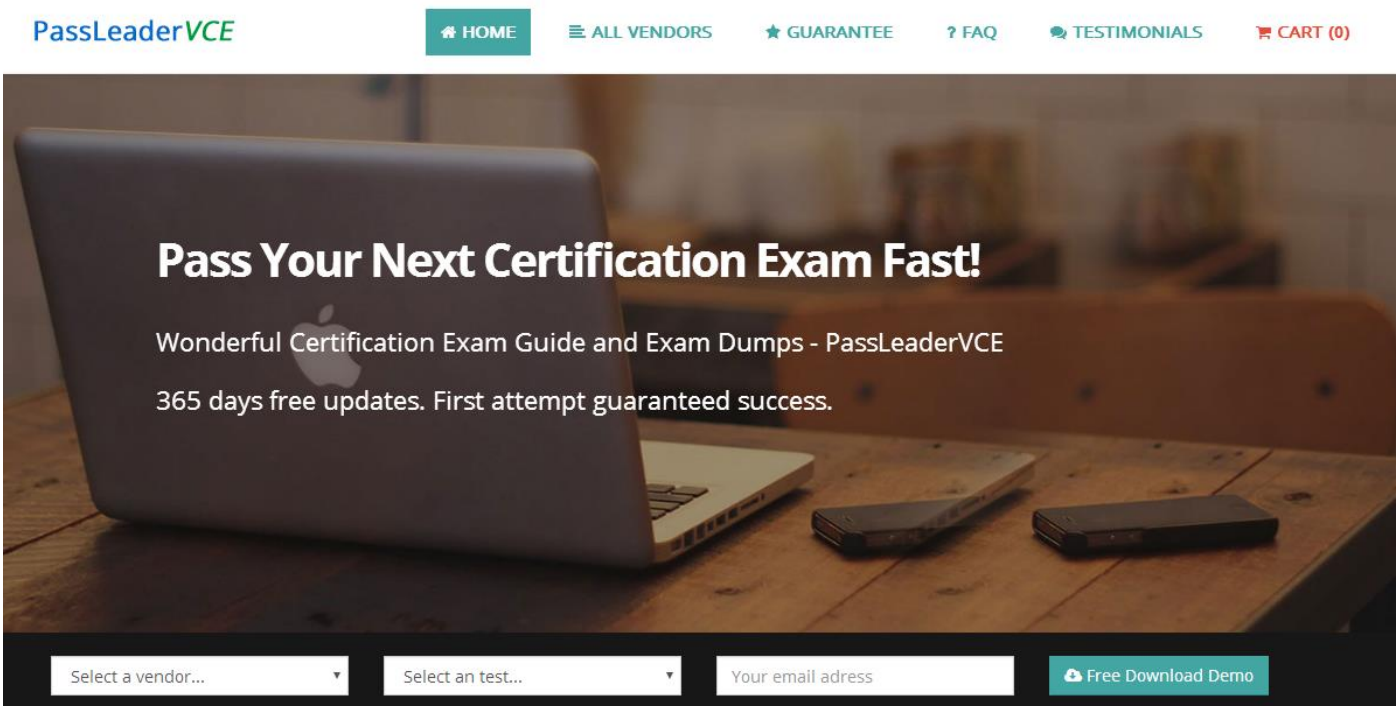

http://www.passleadervce.com/

Wonderful Certification Exam Guide and Exam Dumps- PassLeaderVCE

Exam : **AWS-Solutions-Architect-Professional**

Title : AWS Certified Solutions Architect - Professional

Vendor : Amazon

Version : DEMO

NO.1 A company has an application that uses an Amazon Aurora PostgreSQL DB cluster for the application's database. The DB cluster contains one small primary instance and three larger replica instances. The application runs on an AWS Lambda function. The application makes many short-lived connections to the database's replica instances to perform read-only operations.
During periods of high traffic, the application becomes unreliable and the database reports that too many connections are being established. The frequency of high-traffic periods is unpredictable.
Which solution will improve the reliability of the application?

A. Use Amazon RDS Proxy to create a proxy for the DB cluster. Configure a read-only endpoint for the proxy. Update the Lambda function to connect to the proxy endpoint.

B. Increase the max_connections setting on the DB cluster's parameter group. Reboot all the instances in the DB cluster. Update the Lambda function to connect to the DB cluster endpoint.

C. Configure instance scaling for the DB cluster to occur when the DatabaseConnections metric is close to the max _ connections setting. Update the Lambda function to connect to the Aurora reader endpoint.

D. Use Amazon RDS Proxy to create a proxy for the DB cluster. Configure a read-only endpoint for the Aurora Data API on the proxy. Update the Lambda function to connect to the proxy endpoint.

Answer: A

NO.2 A company needs to build a disaster recovery (DR) solution for its ecommerce website. The web application is hosted on a fleet of t3.large Amazon EC2 instances and uses an Amazon RDS for MySQL DB instance. The EC2 instances are in an Auto Scaling group that extends across multiple Availability Zones.
In the event of a disaster, the web application must fail over to the secondary environment with an RPO of 30 seconds and an R TO of 10 minutes.
Which solution will meet these requirements MOST cost-effectively?

A. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. Create a cross-Region read replica for the DB instance. Set up a backup plan in AWS Backup to create cross-Region backups for the EC2 instances and the DB instance. Create a cron expression to back up the EC2 instances and the DB instance every 30 seconds to the DR Region. Recover the EC2 instances from the latest EC2 backup. Use an Amazon Route 53 geolocation routing policy to automatically fail over to the DR Region in the event of a disaster.

B. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. Create a cross-Region read replica for the DB instance. Set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Region. Run the EC2 instances at the minimum capacity in the DR Region Use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster. Increase the desired capacity of the Auto Scaling group.

C. Set up a backup plan in AWS Backup to create cross-Region backups for the EC2 instances and the DB instance. Create a cron expression to back up the EC2 instances and the DB instance every 30 seconds to the DR Region. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region.
Manually restore the backed-up data on new instances. Use an Amazon Route 53 simple routing policy to automatically fail over to the DR Region in the event of a disaster.

D. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. Create an Amazon Aurora global database. Set up AWS Elastic Disaster Recovery to continuously replicate the

EC2 instances to the DR Region. Run the Auto Scaling group of EC2 instances at full capacity in the DR Region. Use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster.

Answer: B

Explanation:

The company should use infrastructure as code (IaC) to provision the new infrastructure in the DR Region.

The company should create a cross-Region read replica for the DB instance. The company should set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Region. The company should run the EC2 instances at the minimum capacity in the DR Region. The company should use an Amazon Route

53 failover routing policy to automatically fail over to the DR Region in the event of a disaster. The company should increase the desired capacity of the Auto Scaling group. This solution will meet the requirements most cost-effectively because AWS Elastic Disaster Recovery (AWS DRS) is a service that minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery. AWS DRS enables RPOs of seconds and RTOs of minutes1.

AWS DRS continuously replicates data from the source servers to a staging area subnet in the DR Region, where it uses low-cost storage and minimal compute resources to maintain ongoing replication. In the event of a disaster, AWS DRS automatically converts the servers to boot and run natively on AWS and launches recovery instances on AWS within minutes2. By using AWS DRS, the company can save costs by removing idle recovery site resources and paying for the full disaster recovery site only when needed. By creating a cross-Region read replica for the DB instance, the company can have a standby copy of its primary database in a different AWS Region3. By using infrastructure as code (IaC), the company can provision the new infrastructure in the DR Region in an automated and consistent way4. By using an Amazon Route 53 failover routing policy, the company can route traffic to a resource that is healthy or to another resource when the first resource becomes unavailable.

The other options are not correct because:

Using AWS Backup to create cross-Region backups for the EC2 instances and the DB instance would not meet the RPO and RTO requirements. AWS Backup is a service that enables you to centralize and automate data protection across AWS services. You can use AWS Backup to back up your application data across AWS services in your account and across accounts. However, AWS Backup does not provide continuous replication or fast recovery; it creates backups at scheduled intervals and requires manual restoration. Creating backups every 30 seconds would also incur high costs and network bandwidth.

Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not help with disaster recovery. The Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client. It is useful for building applications that interact with Amazon Redshift, but not for replicating or recovering data.

Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not help with disaster recovery. AWS Data Exchange is a service that makes it easy for AWS customers to exchange data in the cloud. You can use AWS Data Exchange to subscribe to a diverse selection of third-party data products or offer your own data products to other AWS customers. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data. It is useful

for sharing query results and views with other users, but not for replicating or recovering data.
References:
https://aws.amazon.com/disaster-recovery/
https://docs.aws.amazon.com/drs/latest/userguide/what-is-drs.html
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html#USER_ReadRepl.XR
https://aws.amazon.com/cloudformation/
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html
https://aws.amazon.com/backup/
https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html
https://aws.amazon.com/data-exchange/
https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html

NO.3 A company has migrated a legacy application to the AWS Cloud. The application runs on three Amazon EC2 instances that are spread across three Availability Zones. One EC2 instance is in each Availability Zone. The EC2 instances are running in three private subnets of the VPC and are set up as targets for an Application Load Balancer (ALB) that is associated with three public subnets.
The application needs to communicate with on-premises systems. Only traffic from IP addresses in the company's IP address range are allowed to access the on-premises systems. The company's security team is bringing only one IP address from its internal IP address range to the cloud. The company has added this IP address to the allow list for the company firewall. The company also has created an Elastic IP address for this IP address.
A solutions architect needs to create a solution that gives the application the ability to communicate with the on-premises systems. The solution also must be able to mitigate failures automatically.
Which solution will meet these requirements?

A. Deploy three NAT gateways, one in each public subnet. Assign the Elastic IP address to the NAT gateways. Turn on health checks for the NAT gateways. If a NAT gateway fails a health check, recreate the NAT gateway and assign the Elastic IP address to the new NAT gateway.

B. Replace the ALB with a Network Load Balancer (NLB). Assign the Elastic IP address to the NLB Turn on health checks for the NLB. In the case of a failed health check, redeploy the NLB in different subnets.

C. Deploy a single NAT gateway in a public subnet. Assign the Elastic IP address to the NAT gateway. Use Amazon CloudWatch with a custom metric to
monitor the NAT gateway. If the NAT gateway is unhealthy, invoke an AWS Lambda function to create a new NAT gateway in a different subnet. Assign the Elastic IP address to the new NAT gateway.

D. Assign the Elastic IP address to the ALB. Create an Amazon Route 53 simple record with the Elastic IP address as the value. Create a Route 53 health check. In the case of a failed health check, recreate the ALB in different subnets.

Answer: C
Explanation:
to connect out from the private subnet you need an NAT gateway and since only one Elastic IP whitelisted on firewall its one NATGateway at time and if AZ failure happens Lambda creates a new NATGATEWAY in a different AZ using the Same Elastic IP ,dont be tempted to select D since application that needs to connect is on a private subnet whose outbound connections use the

NATGateway Elastic IP

NO.4 A finance company hosts a data lake in Amazon S3. The company receives financial data records over SFTP each night from several third parties. The company runs its own SFTP server on an Amazon EC2 instance in a public subnet of a VPC. After the files ate uploaded, they are moved to the data lake by a cron job that runs on the same instance. The SFTP server is reachable on DNS sftp.examWe.com through the use of Amazon Route
53.
What should a solutions architect do to improve the reliability and scalability of the SFTP solution?
A. Move the EC2 instance into an Auto Scaling group. Place the EC2 instance behind an Application Load Balancer (ALB). Update the DNS record sftp.example.com in Route 53 to point to the ALB.
B. Migrate the SFTP server to AWS Transfer for SFTP. Update the DNS record sftp.example.com in Route
53 to point to the server endpoint hostname.
C. Migrate the SFTP server to a file gateway in AWS Storage Gateway. Update the DNS record sflp.example.com in Route 53 to point to the file gateway endpoint.
D. Place the EC2 instance behind a Network Load Balancer (NLB). Update the DNS record sftp.example.com in Route 53 to point to the NLB.
Answer: B
Explanation:
https://aws.amazon.com/aws-transfer-family/faqs/
https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html
https://aws.amazon.com/about-aws/whats-new/2018/11/aws-transfer-for-sftp-fully-managed-sftp-for-s3/?nc1=h_

NO.5 A large company runs workloads in VPCs that are deployed across hundreds of AWS accounts. Each VPC consists to public subnets and private subnets that span across multiple Availability Zones. NAT gateways are deployed in the public subnets and allow outbound connectivity to the internet from the private subnets.
A solutions architect is working on a hub-and-spoke design. All private subnets in the spoke VPCs must route traffic to the internet through an egress VPC. The solutions architect already has deployed a NAT gateway in an egress VPC in a central AWS account.
Which set of additional steps should the solutions architect take to meet these requirements?
A. Create peering connections between the egress VPC and the spoke VPCs. Configure the required routing to allow access to the internet.
B. Create a transit gateway, and share it with the existing AWS accounts. Attach existing VPCs to the transit gateway Configure the required routing to allow access to the internet.
C. Create a transit gateway in every account. Attach the NAT gateway to the transit gateways. Configure the required routing to allow access to the internet.
D. Create an AWS PrivateLink connection between the egress VPC and the spoke VPCs. Configure the required routing to allow access to the internet
Answer: B
Explanation:
https://d1.awsstatic.com/architecture-diagrams/ArchitectureDiagrams/NAT-gateway-centralized-egress-ra.pdf?d

NO.6 A company is planning to migrate an Amazon RDS for Oracle database to an RDS for PostgreSQL DB instance in another AWS account. A solutions architect needs to design a migration strategy that will require no downtime and that will minimize the amount of time necessary to complete the migration. The migration strategy must replicate all existing data and any new data that is created during the migration The target database must be identical to the source database at completion of the migration process All applications currently use an Amazon Route 53 CNAME record as their endpoint for communication with the RDS for Oracle DB instance The RDS for Oracle DB instance is in a private subnet.
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE)

A. Create a new RDS for PostgreSQL DB instance in the target account Use the AWS Schema Conversion Tool (AWS SCT) to migrate the database schema from the source database to the target database

B. Use the AWS Schema Conversion Tool (AWS SCT) to create a new RDS for PostgreSQL DB instance in the target account with the schema and initial data from the source database

C. Configure VPC peering between the VPCs in the two AWS accounts to provide connectivity to both DB instances from the target account. Configure the security groups that are attached to each DB instance to allow traffic on the database port from the VPC in the target account.

D. Temporarily allow the source DB instance to be publicly accessible to provide connectivity from the VPC in the target account Configure the security groups that are attached to each DB instance to allow traffic on the database port from the VPC in the target account.

E. Use AWS Database Migration Service (AWS DMS) in the target account to perform a full load plus change data capture (CDC) migration from the source database to the target database When the migration is complete, change the CNAME record to point to the target DB instance endpoint

F. Use AWS Database Migration Service (AWS DMS) in the target account to perform a change data capture (CDC) migration from the source database to the target database When the migration is complete change the CNAME record to point to the target DB instance endpoint.

Answer: A C E

NO.7 A company has AWS accounts that are in an organization in AWS Organizations. The company wants to track Amazon EC2 usage as a metric. The company's architecture team must receive a daily alert if the EC2 usage is more than 10% higher than the average EC2 usage from the last 30 days.
Which solution will meet these requirements?

A. Configure AWS Budgets in the organization's management account. Specify a usage type of EC2 running hours. Specify a daily period. Set the budget amount to be 10% more than the reported average usage for the last 30 days from AWS Cost Explorer. Configure an alert to notify the architecture team if the usage threshold is met.

B. Configure AWS Cost Anomaly Detection in the organization's management account. Configure a monitor type of AWS Service. Apply a filter of Amazon EC2. Configure an alert subscription to notify the architecture team if the usage is 10% more than the average usage for the last 30 days.

C. Enable AWS Trusted Advisor in the organization's management account. Configure a cost optimization advisory alert to notify the architecture team if the EC2 usage is 10% more than the reported average usage for the last 30 days.

D. Configure Amazon Detective in the organization's management account. Configure an EC2 usage anomaly alert to notify the architecture team if Detective identifies a usage anomaly of more than 10%.

Answer: B

Explanation:

AWS Cost Anomaly Detection is a feature of the AWS Cost Management suite that leverages machine learning to enable continuous monitoring of your AWS costs and usage, allowing you to identify unexpected and abnormal spending1. You can create cost monitors that evaluate specific AWS services, member accounts, cost allocation tags, or cost categories based on your AWS account structure2. You can also configure alert subscriptions that notify you when a cost monitor detects an anomaly that meets your threshold2. In this case, you can create a cost monitor with a monitor type of AWS Service and apply a filter of Amazon EC2 to track the EC2 usage as a metric. You can then configure an alert subscription to notify the architecture team if the usage is 10% more than the average usage for the last 30 days, which is the anomaly detection period used by AWS Cost Anomaly Detection3.

NO.8 A company runs many workloads on AWS and uses AWS Organizations to manage its accounts. The workloads are hosted on Amazon EC2. AWS Fargate. and AWS Lambda. Some of the workloads have unpredictable demand. Accounts record high usage in some months and low usage in other months.

The company wants to optimize its compute costs over the next 3 years A solutions architect obtains a

6-month average for each of the accounts across the organization to calculate usage.

Which solution will provide the MOST cost savings for all the organization's compute usage?

A. Purchase Reserved Instances for the organization to match the size and number of the most common EC2 instances from the member accounts.

B. Purchase a Compute Savings Plan for the organization from the management account by using the recommendation at the management account level

C. Purchase Reserved Instances for each member account that had high EC2 usage according to the data from the last 6 months.

D. Purchase an EC2 Instance Savings Plan for each member account from the management account based on EC2 usage data from the last 6 months.

Answer: B

NO.9 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:

VPCCIDR 10 0 0 0/23

AZ1 subnet CIDR: 10 0 0 0724

AZ2 subnet CIDR: 10.0.1 0724

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 Delete and re-create the AZ1 subnet using half the previous address space Adjust the Auto Scaling group to also use the new AZI subnet When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only Remove the current AZ2 subnet Create a new AZ2 subnet using the second half of the address space from the original AZ1 subnet 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.

B. Terminate the EC2 instances in the AZ1 subnet Delete and re-create the AZ1 subnet using hall the address space. Update the Auto Scaling group to use this new subnet. Repeat this for the second AZ. Define a new subnet in AZ3: then update the Auto Scaling group to target all three new subnets

C. Create a new VPC with the same IPv4 address space and define three subnets, with one for each AZ Update the existing Auto Scaling group to target the new subnets in the new VPC

D. Update the Auto Scaling group to use the AZ2 subnet only Update the AZ1 subnet to have halt the previous address space Adjust the Auto Scaling group to also use the AZ1 subnet again. When the instances are healthy, adjust the Auto Seating group to use the AZ1 subnet only. Update the current AZ2 subnet and assign the second half of the address space from the original AZ1 subnet 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://repost.aws/knowledge-center/vpc-ip-address-range

NO.10 A company hosts a software as a service (SaaS) solution on AWS. The solution has an Amazon API Gateway API that serves an HTTPS endpoint. The API uses AWS Lambda functions for compute. The Lambda functions store data in an Amazon Aurora Serverless VI database.

The company used the AWS Serverless Application Model (AWS SAM) to deploy the solution. The solution extends across multiple Availability Zones and has no disaster recovery (DR) plan.

A solutions architect must design a DR strategy that can recover the solution in another AWS Region. The solution has an R TO of 5 minutes and an RPO of 1 minute.

What should the solutions architect do to meet these requirements?

A. Create a read replica of the Aurora Serverless VI database in the target Region. Use AWS SAM to create a runbook to deploy the solution to the target Region. Promote the read replica to primary in case of disaster.

B. Change the Aurora Serverless VI database to a standard Aurora MySQL global database that extends across the source Region and the target Region. Use AWS SAM to create a runbook to deploy the solution to the target Region.

C. Create an Aurora Serverless VI DB cluster that has multiple writer instances in the target Region. Launch the solution in the target Region. Configure the two Regional solutions to work in an active-passive configuration.

D. Change the Aurora Serverless VI database to a standard Aurora MySQL global database that extends across the source Region and the target Region. Launch the solution in the target Region. Configure the two Regional solutions to work in an active-passive configuration.

Answer: D

Explanation: This option allows the solutions architect to use Aurora global database to replicate data across multiple AWS Regions with low latency and high availability1. By launching the solution in the target Region, the solutions architect can ensure that the API Gateway, Lambda functions, and other

resources are ready to serve traffic in case of a disaster in the source Region. By configuring the two Regional solutions to work in an active-passive configuration, the solutions architect can minimize costs and avoid data conflicts by having only one primary Region that accepts write operations and one secondary Region that serves as a standby2. The RTO and RPO requirements can be met by using Aurora global database, which supports sub-second failover times and near real-time replication1.
References:
Working with Amazon Aurora global database
Active-passive failover

NO.11 A company is using an organization in AWS Organizations to manage hundreds of AWS accounts. A solutions architect is working on a solution to provide baseline protection for the Open Web Application Security Project (OWASP) top 10 web application vulnerabilities. The solutions architect is using AWS WAF for all existing and new Amazon CloudFront distributions that are deployed within the organization.
Which combination of steps should the solutions architect take to provide the baseline protection? (Select THREE.)

A. Enable AWS Config in all accounts.

B. Enable Amazon GuardDuty in all accounts.

C. Enable all features for the organization.

D. Use AWS Firewall Manager to deploy AWS WAF rules in all accounts for all CloudFront distributions.

E. Use AWS Shield Advanced to deploy AWS WAF rules in all accounts for all CloudFront distributions.

F. Use AWS Security Hub to deploy AWS WAF rules in all accounts for all CloudFront distributions.

Answer: C D E
Explanation:
Enabling all features for the organization will enable using AWS Firewall Manager to centrally configure and manage firewall rules across multiple AWS accounts1. Using AWS Firewall Manager to deploy AWS WAF rules in all accounts for all CloudFront distributions will enable providing baseline protection for the OWASP top 10 web application vulnerabilities2. AWS Firewall Manager supports AWS WAF rules that can help protect against common web exploits such as SQL injection and cross-site scripting3. Configuring redirection of HTTP requests to HTTPS requests in CloudFront will enable encrypting the data in transit using SSL/TLS.

NO.12 A company has an organization that has many AWS accounts in AWS Organizations. A solutions architect must improve how the company manages common security group rules for the AWS accounts in the organization.
The company has a common set of IP CIDR ranges in an allow list in each AWS account to allow access to and from the company's on-premises network.
Developers within each account are responsible for adding new IP CIDR ranges to their security groups. The security team has its own AWS account. Currently, the security team notifies the owners of the other AWS accounts when changes are made to the allow list.
The solutions architect must design a solution that distributes the common set of CIDR ranges across all accounts.
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Set up an Amazon Simple Notification Service (Amazon SNS) topic in the security team's AWS

account. Deploy an AWS Lambda function in each AWS account. Configure the Lambda function to run every time an SNS topic receives a message. Configure the Lambda function to take an IP address as input and add it to a list of security groups in the account. Instruct the security team to distribute changes by publishing messages to its SNS topic.

B. Create new customer-managed prefix lists in each AWS account within the organization. Populate the prefix lists in each account with all internal CIDR ranges. Notify the owner of each AWS account to allow the new customer-managed prefix list IDs in their accounts in their security groups. Instruct the security team to share updates with each AWS account owner.

C. Create a new customer-managed prefix list in the security team's AWS account. Populate the customer-managed prefix list with all internal CIDR ranges. Share the customer-managed prefix list with the organization by using AWS Resource Access Manager. Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups.

D. Create an IAM role in each account in the organization. Grant permissions to update security groups.

Deploy an AWS Lambda function in the security team's AWS account. Configure the Lambda function to take a list of internal IP addresses as input, assume a role in each organization account, and add the list of IP addresses to the security groups in each account.

Answer: C

Explanation:

Create a new customer-managed prefix list in the security team's AWS account. Populate the customer-managed prefix list with all internal CIDR ranges. Share the customer-managed prefix list with the organization by using AWS Resource Access Manager. Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups. This solution meets the requirements with the least amount of operational overhead as it requires the security team to create and maintain a single customer-managed prefix list, and share it with the organization using AWS Resource Access Manager. The owners of each AWS account are then responsible for allowing the prefix list in their security groups, which eliminates the need for the security team to manually notify each account owner when changes are made. This solution also eliminates the need for a separate AWS Lambda function in each account, reducing the overall complexity of the solution.

NO.13 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

NO.14 A company needs to migrate its customer transactions database from on premises to AWS. The database resides on an Oracle DB instance that runs on a Linux server. According to a new security requirement, the company must rotate the database password each year.
Which solution will meet these requirements with the LEAST operational overhead?
A. Convert the database to Amazon DynamoDB by using the AWS Schema Conversion Tool (AWS SCT).
Store the password in AWS Systems Manager Parameter Store. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function for yearly password rotation.
B. Migrate the database to Amazon RDS for Oracle. Store the password in AWS Secrets Manager. Turn on automatic rotation. Configure a yearly rotation schedule.
C. Migrate the database to an Amazon EC2 instance. Use AWS Systems Manager Parameter Store to keep and rotate the connection string by using an AWS Lambda function on a yearly schedule
D. Migrate the database to Amazon Neptune by using the AWS Schema Conversion Tool {AWS SCT). Create an Amazon CloudWatch alarm to invoke an AWS Lambda function for yearly password rotation.
Answer: B

NO.15 A company is developing and hosting several projects in the AWS Cloud. The projects are developed across multiple AWS accounts under the same organization in AWS Organizations. The company requires the cost lor cloud infrastructure to be allocated to the owning project. The team responsible for all of the AWS accounts has discovered that several Amazon EC2 instances are lacking the Project tag used for cost allocation.
Which actions should a solutions architect take to resolve the problem and prevent it from happening in the future? (Select THREE.)
A. Create an AWS Config rule in each account to find resources with missing tags.
B. Create an SCP in the organization with a deny action for ec2:RunInstances if the Project tag is missing.
C. Use Amazon Inspector in the organization to find resources with missing tags.
D. Create an IAM policy in each account with a deny action for ec2:RunInstances if the Project tag is missing.
E. Create an AWS Config aggregator for the organization to collect a list of EC2 instances with the missing Project tag.
F. Use AWS Security Hub to aggregate a list of EC2 instances with the missing Project tag.
Answer: A B E
Explanation:
https://docs.aws.amazon.com/config/latest/developerguide/config-rule-multi-account-deployment.html
https://docs.aws.amazon.com/config/latest/developerguide/aggregate-data.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_tagging.html

NO.16 A company is running an application in the AWS Cloud. The core business logic is running on a

set of Amazon EC2 instances in an Auto Scaling group. An Application Load Balancer (ALB) distributes traffic to the EC2 instances. Amazon Route 53 record api.example.com is pointing to the ALB.
The company's development team makes major updates to the business logic. The company has a rule that when changes are deployed, only 10% of customers can receive the new logic during a testing window. A customer must use the same version of the business logic during the testing window.
How should the company deploy the updates to meet these requirements?

A. Create a second ALB, and deploy the new logic to a set of EC2 instances in a new Auto Scaling group.
Configure the ALB to distribute traffic to the EC2 instances. Update the Route 53 record to use weighted routing, and point the record to both of the ALBs.

B. Create a second target group that is referenced by the ALB. Deploy the new logic to EC2 instances in this new target group. Update the ALB listener rule to use weighted target groups. Configure ALB target group stickiness.

C. Create a new launch configuration for the Auto Scaling group. Specify the launch configuration to use the AutoScalingRollingUpdate policy, and set the MaxBatchSize option to 10. Replace the launch configuration on the Auto Scaling group. Deploy the changes.

D. Create a second Auto Scaling group that is referenced by the ALB. Deploy the new logic on a set of EC2 instances in this new Auto Scaling group. Change the ALB routing algorithm to least outstanding requests (LOR). Configure ALB session stickiness.

Answer: B
Explanation:
The company should create a second target group that is referenced by the ALB. The company should deploy the new logic to EC2 instances in this new target group. The company should update the ALB listener rule to use weighted target groups. The company should configure ALB target group stickiness. This solution will meet the requirements because weighted target groups are a feature that enables you to distribute traffic across multiple target groups using a single listener rule. You can specify a weight for each target group, which determines the percentage of requests that are routed to that target group. For example, if you specify two target groups, each with a weight of 10, each target group receives half the requests1. By creating a second target group and deploying the new logic to EC2 instances in this new target group, the company can have two versions of its business logic running in parallel. By updating the ALB listener rule to use weighted target groups, the company can control how much traffic is sent to each version. By configuring ALB target group stickiness, the company can ensure that a customer uses the same version of the business logic during the testing window. Target group stickiness is a feature that enables you to bind a user's session to a specific target within a target group for the duration of the session2.
The other options are not correct because:
Creating a second ALB and deploying the new logic to a set of EC2 instances in a new Auto Scaling group would not be as cost-effective or simple as using weighted target groups. A second ALB would incur additional charges and require more configuration and management. Updating the Route 53 record to use weighted routing would not ensure that a customer uses the same version of the business logic during the testing window, as DNS caching could affect how requests are routed.
Creating a new launch configuration for the Auto Scaling group and replacing it on the Auto Scaling group would not allow for gradual traffic shifting between versions. A launch configuration is a template that an Auto Scaling group uses to launch EC2 instances. You can specify information such

as the AMI ID, instance type, key pair, security groups, and block device mapping for your instances3. However, replacing the launch configuration on an Auto Scaling group would affect all instances in that group, not just 10% of customers.
Creating a second Auto Scaling group and changing the ALB routing algorithm to least outstanding requests (LOR) would not allow for controlled traffic shifting between versions. A second Auto Scaling group would require more configuration and management. The LOR routing algorithm is a feature that enables you to route traffic based on how quickly targets respond to requests. The load balancer selects a target from the target group with the fewest outstanding requests4. However, this algorithm does not take into account customer sessions or weights.
References:
https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-listeners.html#listener-
https://docs.aws.amazon.com/elasticloadbalancing/latest/application/sticky-sessions.html
https://docs.aws.amazon.com/autoscaling/ec2/userguide/LaunchConfiguration.html
https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-target-groups.html#rou

NO.17 A company has deployed its database on an Amazon RDS for MySQL DB instance in the us-east-1 Region.
The company needs to make its data available to customers in Europe. The customers in Europe must have access to the same data as customers in the United States (US) and will not tolerate high application latency or stale data. The customers in Europe and the customers in the US need to write to the database. Both groups of customers need to see updates from the other group in real time.
Which solution will meet these requirements?
A. Create an Amazon Aurora MySQL replica of the RDS for MySQL DB instance. Pause application writes to the RDS DB instance. Promote the Aurora Replica to a standalone DB cluster. Reconfigure the application to use the Aurora database and resume writes. Add eu-west-1 as a secondary Region to the
06 cluster. Enable write forwarding on the DB cluster. Deploy the application in eu-west-1. Configure the application to use the Aurora MySQL endpoint in eu- west-1.
B. Add a cross-Region replica in eu-west-1 for the RDS for MySQL DB instance. Configure the replica to replicate write queries back to the primary DB instance. Deploy the application in eu-west-1. Configure the application to use the RDS for MySQL endpoint in eu-west-1.
C. Copy the most recent snapshot from the RDS for MySQL DB instance to eu-west-1. Create a new RDS for MySQL DB instance in eu-west-1 from the snapshot. Configure MySQL logical replication from us-east-1 to eu-west-1. Enable write forwarding on the DB cluster. Deploy the application in eu-west-1.
Configure the application to use the RDS for MySQL endpoint in eu-west-1.
D. Convert the RDS for MySQL DB instance to an Amazon Aurora MySQL DB cluster. Add eu-west-1 as a secondary Region to the DB cluster. Enable write forwarding on the DB cluster. Deploy the application in eu-west-1. Configure the application to use the Aurora MySQL endpoint in eu-west-1.
Answer: D
Explanation:
The company should use AWS Amplify to create a static website for uploads of media files. The company should use Amplify Hosting to serve the website through Amazon CloudFront. The company

should use Amazon S3 to store the uploaded media files. The company should use Amazon Cognito to authenticate users.
This solution will meet the requirements with the least operational overhead because AWS Amplify is a complete solution that lets frontend web and mobile developers easily build, ship, and host full-stack applications on AWS, with the flexibility to leverage the breadth of AWS services as use cases evolve. No cloud expertise needed1. By using AWS Amplify, the company can refactor the application to a serverless architecture that reduces operational complexity and costs. AWS Amplify offers the following features and benefits:
Amplify Studio: A visual interface that enables you to build and deploy a full-stack app quickly, including frontend UI and backend.
Amplify CLI: A local toolchain that enables you to configure and manage an app backend with just a few commands.
Amplify Libraries: Open-source client libraries that enable you to build cloud-powered mobile and web apps.
Amplify UI Components: Open-source design system with cloud-connected components for building feature-rich apps fast.
Amplify Hosting: Fully managed CI/CD and hosting for fast, secure, and reliable static and server-side rendered apps.
By using AWS Amplify to create a static website for uploads of media files, the company can leverage Amplify Studio to visually build a pixel-perfect UI and connect it to a cloud backend in clicks. By using Amplify Hosting to serve the website through Amazon CloudFront, the company can easily deploy its web app or website to the fast, secure, and reliable AWS content delivery network (CDN), with hundreds of points of presence globally. By using Amazon S3 to store the uploaded media files, the company can benefit from a highly scalable, durable, and cost-effective object storage service that can handle any amount of data2. By using Amazon Cognito to authenticate users, the company can add user sign-up, sign-in, and access control to its web app with a fully managed service that scales to support millions of users3.
The other options are not correct because:
Using AWS Application Migration Service to migrate the application server to Amazon EC2 instances would not refactor the application or accelerate development. AWS Application Migration Service (AWS MGN) is a service that enables you to migrate physical servers, virtual machines (VMs), or cloud servers from any source infrastructure to AWS without requiring agents or specialized tools. However, this would not address the challenges of overutilization and data uploads failures. It would also not reduce operational overhead or costs compared to a serverless architecture.
Creating a static website for uploads of media files and using AWS AppSync to create an API would not be as simple or fast as using AWS Amplify. AWS AppSync is a service that enables you to create flexible APIs for securely accessing, manipulating, and combining data from one or more data sources.
However, this would require more configuration and management than using Amplify Studio and Amplify Hosting. It would also not provide authentication features like Amazon Cognito.
Setting up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the application would not be as suitable as using Amazon Cognito. AWS Single Sign-On (AWS SSO) is a service that enables you to centrally manage SSO access and user permissions across multiple AWS accounts and business applications. However, this service is designed for enterprise customers who need to manage access for employees or partners across multiple resources. It is not intended for authenticating end users of web or mobile apps.

References:
https://aws.amazon.com/amplify/
https://aws.amazon.com/s3/
https://aws.amazon.com/cognito/
https://aws.amazon.com/mgn/
https://aws.amazon.com/appsync/
https://aws.amazon.com/single-sign-on/

NO.18 A company gives users the ability to upload images from a custom application. The upload process invokes an AWS Lambda function that processes and stores the image in an Amazon S3 bucket. The application invokes the Lambda function by using a specific function version ARN.
The Lambda function accepts image processing parameters by using environment variables. The company often adjusts the environment variables of the Lambda function to achieve optimal image processing output.
The company tests different parameters and publishes a new function version with the updated environment variables after validating results. This update process also requires frequent changes to the custom application to invoke the new function version ARN. These changes cause interruptions for users.
A solutions architect needs to simplify this process to minimize disruption to users.
Which solution will meet these requirements with the LEAST operational overhead?

A. Directly modify the environment variables of the published Lambda function version. Use the $LATEST version to test image processing parameters.

B. Create an Amazon DynamoDB table to store the image processing parameters. Modify the Lambda function to retrieve the image processing parameters from the DynamoDB table.

C. Directly code the image processing parameters within the Lambda function and remove the environment variables. Publish a new function version when the company updates the parameters.

D. Create a Lambda function alias. Modify the client application to use the function alias ARN. Reconfigure the Lambda alias to point to new versions of the function when the company finishes testing.

Answer: D

Explanation:
A Lambda function alias allows you to point to a specific version of a function and also can be updated to point to a new version of the function without modifying the client application. This way, the company can test different versions of the function with different environment variables and, once the optimal parameters are found, update the alias to point to the new version, without the need to update the client application.
By using this approach, the company can simplify the process of updating the environment variables, minimize disruption to users, and reduce the operational overhead.
Reference:
AWS Lambda documentation: https://aws.amazon.com/lambda/
AWS Lambda Aliases documentation: https://docs.aws.amazon.com/lambda/latest/dg/aliases-intro.html AWS Lambda versioning and aliases documentation:
https://aws.amazon.com/blogs/compute/versioning-aliases-in-aws-lambda/

NO.19 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. Include the necessary IAM roles for each administrator.

B. Create the OrganizationAccountAccessPolicy IAM policy in each member account. Connect the member accounts to the management account by using cross- account access.

C. Create the OrganizationAccountAccessRole IAM role in each member account. Grant permission to the management account to assume the IAM role.

D. Create the OrganizationAccountAccessRole IAM role in the management account. Attach the AdministratorAccess AWS managed policy to the IAM role. Assign the IAM role to the administrators in each member account.

Answer: C

NO.20 A company is deploying a distributed in-memory database on a fleet of Amazon EC2 instances. The fleet consists of a primary node and eight worker nodes. The primary node is responsible for monitoring cluster health, accepting user requests, distributing user requests to worker nodes, and sending an aggregate response back to a client. Worker nodes communicate with each other to replicate data partitions.
The company requires the lowest possible networking latency to achieve maximum performance.
Which solution will meet these requirements?

A. Launch memory optimized EC2 instances in a partition placement group.

B. Launch compute optimized EC2 instances in a partition placement group.

C. Launch memory optimized EC2 instances in a cluster placement group

D. Launch compute optimized EC2 instances in a spread placement group.

Answer: C

NO.21 A software company needs to create short-lived test environments to test pull requests as part of its development process. Each test environment consists of a single Amazon EC2 instance that is in an Auto Scaling group.
The test environments must be able to communicate with a central server to report test results. The central server is located in an on-premises data center. A solutions architect must implement a solution so that the company can create and delete test environments without any manual intervention. The company has created a transit gateway with a VPN attachment to the on-premises network.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an AWS CloudFormation template that contains a transit gateway attachment and related routing configurations. Create a CloudFormation stack set that includes this template. Use CloudFormation StackSets to deploy a new stack for each VPC in the account. Deploy a new VPC for each test environment.

B. Create a single VPC for the test environments. Include a transit gateway attachment and related routing configurations. Use AWS CloudFormation to deploy all test environments into the VPC.

C. Create a new OU in AWS Organizations for testing. Create an AWS CloudFormation template that

contains a VPC, necessary networking resources, a transit gateway attachment, and related routing configurations. Create a CloudFormation stack set that includes this template. Use CloudFormation StackSets for deployments into each account under the testing 01.1. Create a new account for each test environment.

D. Convert the test environment EC2 instances into Docker images. Use AWS CloudFormation to configure an Amazon Elastic Kubernetes Service (Amazon EKS) cluster in a new VPC, create a transit gateway attachment, and create related routing configurations. Use Kubernetes to manage the deployment and lifecycle of the test environments.

Answer: B

Explanation: This option allows the company to use a single VPC to host multiple test environments that are isolated from each other by using different subnets and security groups1. By including a transit gateway attachment and related routing configurations, the company can enable the test environments to communicate with the central server in the on-premises data center through a VPN connection2. By using AWS CloudFormation to deploy all test environments into the VPC, the company can automate the creation and deletion of test environments without any manual intervention3. This option also minimizes the operational overhead by reducing the number of VPCs, accounts, and resources that need to be managed.

References:

Working with VPCs and subnets

Working with transit gateways

Working with AWS CloudFormation stacks

NO.22 A company has applications in an AWS account that is named Source. The account is in an organization in AWS Organizations. One of the applications uses AWS Lambda functions and store's inventory data in an Amazon Aurora database. The application deploys the Lambda functions by using a deployment package. The company has configured automated backups for Aurora.

The company wants to migrate the Lambda functions and the Aurora database to a new AWS account that is named Target. The application processes critical data, so the company must minimize downtime.

Which solution will meet these requirements?

A. Download the Lambda function deployment package from the Source account. Use the deployment package and create new Lambda functions in the Target account. Share the automated Aurora DB cluster snapshot with the Target account.

B. Download the Lambda function deployment package from the Source account. Use the deployment package and create new Lambda functions in the Target account Share the Aurora DB cluster with the Target account by using AWS Resource Access Manager {AWS RAM). Grant the Target account permission to clone the Aurora DB cluster.

C. Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions and the Aurora DB cluster with the Target account. Grant the Target account permission to clone the Aurora DB cluster.

D. Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions with the Target account. Share the automated Aurora DB cluster snapshot with the Target account.

Answer: C

Explanation:

This solution uses a combination of AWS Resource Access Manager (RAM) and automated backups to migrate the Lambda functions and the Aurora database to the Target account while minimizing

downtime. In this solution, the Lambda function deployment package is downloaded from the Source account and used to create new Lambda functions in the Target account. The Aurora DB cluster is shared with the Target account using AWS RAM and the Target account is granted permission to clone the Aurora DB cluster, allowing for a new copy of the Aurora database to be created in the Target account. This approach allows for the data to be migrated to the Target account while minimizing downtime, as the Target account can use the cloned Aurora database while the original Aurora database continues to be used in the Source account.

NO.23 A company is using AWS CodePipeline for the CI/CD of an application to an Amazon EC2 Auto Scaling group. All AWS resources are defined in AWS CloudFormation templates. The application artifacts are stored in an Amazon S3 bucket and deployed to the Auto Scaling group using instance user data scripts.
As the application has become more complex, recent resource changes in the CloudFormation templates have caused unplanned downtime.
How should a solutions architect improve the CI/CD pipeline to reduce the likelihood that changes in the templates will cause downtime?
A. Adapt the deployment scripts to detect and report CloudFormation error conditions when performing deployments. Write test plans for a testing team to execute in a non-production environment before approving the change for production.
B. Implement automated testing using AWS CodeBuild in a test environment. Use CloudFormation change sets to evaluate changes before deployment. Use AWS CodeDeploy to leverage blue/green deployment patterns to allow evaluations and the ability to revert changes, if needed.
C. Use plugins for the integrated development environment (IDE) to check the templates for errors, and use the AWS CLI to validate that the templates are correct. Adapt the deployment code to check for error conditions and generate notifications on errors. Deploy to a test environment and execute a manual test plan before approving the change for production.
D. Use AWS CodeDeploy and a blue/green deployment pattern with CloudFormation to replace the user data deployment scripts. Have the operators log in to running instances and go through a manual test plan to verify the application is running as expected.
Answer: B

NO.24 A company has a critical application in which the data tier is deployed in a single AWS Region. The data tier uses an Amazon DynamoDB table and an Amazon Aurora MySQL DB cluster. The current Aurora MySQL engine version supports a global database. The application tier is already deployed in two Regions.
Company policy states that critical applications must have application tier components and data tier components deployed across two Regions. The RTO and RPO must be no more than a few minutes each. A solutions architect must recommend a solution to make the data tier compliant with company policy.
Which combination of steps will meet these requirements? (Choose two.)
A. Add another Region to the Aurora MySQL DB cluster
B. Add another Region to each table in the Aurora MySQL DB cluster
C. Set up scheduled cross-Region backups for the DynamoDB table and the Aurora MySQL DB cluster
D. Convert the existing DynamoDB table to a global table by adding another Region to its configuration

E. Use Amazon Route 53 Application Recovery Controller to automate database backup and recovery to the secondary Region

Answer: A D

Explanation:

The company should use Amazon Aurora global database and Amazon DynamoDB global table to deploy the data tier components across two Regions. Amazon Aurora global database is a feature that allows a single Aurora database to span multiple AWS Regions, enabling low-latency global reads and fast recovery from Region-wide outages1. Amazon DynamoDB global table is a feature that allows a single DynamoDB table to span multiple AWS Regions, enabling low-latency global reads and writes and fast recovery from Region-wide outages2.

References:

https://aws.amazon.com/rds/aurora/global-database/

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/globaltables_HowItWorks.html

https://aws.amazon.com/route53/application-recovery-controller/

NO.25 A company is running an application in the AWS Cloud. The application uses AWS Lambda functions and Amazon Elastic Container Service (Amazon ECS) containers that run with AWS Fargate technology as its primary compute. The load on the application is irregular. The application experiences long periods of no usage, followed by sudden and significant increases and decreases in traffic. The application is write-heavy and stores data in an Amazon Aurora MySQL database. The database runs on an Amazon RDS memory optimized DB instance that is not able to handle the load. What is the MOST cost-effective way for the company to handle the sudden and significant changes in traffic?

A. Add additional read replicas to the database. Purchase Instance Savings Plans and RDS Reserved Instances.

B. Migrate the database to an Aurora multi-master DB cluster. Purchase Instance Savings Plans.

C. Migrate the database to an Aurora global database. Purchase Compute Savings Plans and RDS Reserved Instances.

D. Migrate the database to Aurora Serverless v1. Purchase Compute Savings Plans.

Answer: D

NO.26 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 gateway. Delete the existing private virtual interface from the existing connection. Create the second Direct Connect connection. Create a new private virtual interlace on each connection, and connect both private victual interfaces to the Direct Connect gateway. Connect the Direct Connect gateway to the single VPC.

B. Keep the existing private virtual interface. Create the second Direct Connect connection. Create a new private virtual interface on the new connection, and connect the new private virtual interface to the single VPC.
C. Keep the existing private virtual interface. Create the second Direct Connect connection. Create a new public virtual interface on the new connection, and connect the new public virtual interface to the single VPC.
D. Provision a transit gateway. Delete the existing private virtual interface from the existing connection.
Create the second Direct Connect connection. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the transit gateway. 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

NO.27 A company is running an application on Amazon EC2 instances in the AWS Cloud. The application is using a MongoDB database with a replica set as its data tier. The MongoDB database is installed on systems in the company's on-premises data center and is accessible through an AWS Direct Connect connection to the data center environment.
A solutions architect must migrate the on-premises MongoDB database to Amazon DocumentDB (with MongoDB compatibility).
Which strategy should the solutions architect choose to perform this migration?
A. Create a fleet of EC2 instances. Install MongoDB Community Edition on the EC2 instances, and create a database. Configure continuous synchronous replication with the database that is running in the on-premises data center.
B. Create an AWS Database Migration Service (AWS DMS) replication instance. Create a source endpoint for the on-premises MongoDB database by using change data capture (CDC). Create a target endpoint for the Amazon DocumentDB database. Create and run a DMS migration task.
C. Create a data migration pipeline by using AWS Data Pipeline. Define data nodes for the on-premises MongoDB database and the Amazon DocumentDB database. Create a scheduled task to run the data pipeline.
D. Create a source endpoint for the on-premises MongoDB database by using AWS Glue crawlers. Configure continuous asynchronous replication between the MongoDB database and the Amazon DocumentDB database.
Answer: B
Explanation:
https://aws.amazon.com/getting-started/hands-on/move-to-managed/migrate-mongodb-to-documentdb/

NO.28 A solutions architect is determining the DNS strategy for an existing VPC. The VPC is provisioned to use the
10.24.34.0/24 CIDR block. The VPC also uses Amazon Route 53 Resolver for DNS. New requirements

mandate that DNS queries must use private hosted zones. Additionally, instances that have public IP addresses must receive corresponding public hostnames.
Which solution will meet these requirements to ensure that the domain names are correctly resolved within the VPC?

A. Create a private hosted zone. Activate the enableDnsSupport attribute and the enableDnsHostnames attribute for the VPC. Update the VPC DHCP options set to include domain-name-servers-10.24.34.2.

B. Create a private hosted zone. Associate the private hosted zone with the VPC. Activate the enableDnsSupport attribute and the enableDnsHostnames attribute for the VPC. Create a new VPC DHCP options set, and configure domain-name-servers=AmazonProvidedDNS. Associate the new DHCP options set with the VPC.

C. Deactivate the enableDnsSupport attribute for the VPC. Activate the enableDnsHostnames attribute for the VPC. Create a new VPC DHCP options set, and configure domain-name-servers=10.24.34.2.
Associate the new DHCP options set with the VPC.

D. Create a private hosted zone. Associate the private hosted zone with the VPC. Activate the enableDnsSupport attribute for the VPC. Deactivate the enableDnsHostnames attribute for the VPC. Update the VPC DHCP options set to include domain-name-servers=AmazonProvidedDNS.

Answer: B

Explanation: This option allows the solutions architect to use a private hosted zone to host DNS records that are only accessible within the VPC1. By associating the private hosted zone with the VPC, the solutions architect can ensure that DNS queries from the VPC are routed to the private hosted zone2. By activating the enableDnsSupport attribute and the enableDnsHostnames attribute for the VPC, the solutions architect can enable DNS resolution and hostname assignment for instances in the VPC3. By creating a new VPC DHCP options set, and configuring domain-name-servers=AmazonProvidedDNS, the solutions architect can use Amazon-provided DNS servers to resolve DNS queries from instances in the VPC4. By associating the new DHCP options set with the VPC, the solutions architect can apply the DNS settings to all instances in the VPC5.
References:
What is Amazon Route 53 Resolver?
Associating a private hosted zone with your VPC
Using DNS with your VPC
DHCP options sets
Modifying your DHCP options

NO.29 A company is running a workload that consists of thousands of Amazon EC2 instances. The workload is running in a VPC that contains several public subnets and private subnets. The public subnets have a route for
0.0.0.0/0 to an existing internet gateway. The private subnets have a route for 0.0.0.0/0 to an existing NAT gateway.
A solutions architect needs to migrate the entire fleet of EC2 instances to use IPv6. The EC2 instances that are in private subnets must not be accessible from the public internet.
What should the solutions architect do to meet these requirements?

A. Update the existing VPC, and associate a custom IPv6 CIDR block with the VPC and all subnets. Update all the VPC route tables, and add a route for ::/0 to the internet gateway.

B. Update the existing VPC, and associate an Amazon-provided IPv6 CIDR block with the VPC and all subnets. Update the VPC route tables for all private subnets, and add a route for ::/0 to the NAT gateway.
C. Update the existing VPC, and associate an Amazon-provided IPv6 CIDR block with the VPC and all subnets. Create an egress-only internet gateway. Update the VPC route tables for all private subnets, and add a route for ::/0 to the egress-only internet gateway.
D. Update the existing VPC, and associate a custom IPv6 CIDR block with the VPC and all subnets. Create a new NAT gateway, and enable IPv6 support. Update the VPC route tables for all private subnets, and add a route for ::/0 to the IPv6-enabled NAT gateway.
Answer: C

NO.30 A company is running a critical application that uses an Amazon RDS for MySQL database to store data. The RDS DB instance is deployed in Multi-AZ mode.
A recent RDS database failover test caused a 40-second outage to the application A solutions architect needs to design a solution to reduce the outage time to less than 20 seconds.
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)
A. Use Amazon ElastiCache for Memcached in front of the database
B. Use Amazon ElastiCache for Redis in front of the database.
C. Use RDS Proxy in front of the database
D. Migrate the database to Amazon Aurora MySQL
E. Create an Amazon Aurora Replica
F. Create an RDS for MySQL read replica
Answer: C D E
Explanation:
Migrate the database to Amazon Aurora MySQL. - Create an Amazon Aurora Replica. - Use RDS Prox y in front of the database. - These options are correct because they address the requirement of reducing the failover time to less than 20 seconds. Migrating to Amazon Aurora MySQL and creating an Aurora replica can reduce the failover time to less than 20 seconds. Aurora has a built-in, fault-tolerant storage system that can automatically detect and repair failures. Additionally, Aurora has a feature called "Aurora Global Database" which allows you to create read-only replicas across multiple AWS regions which can further help to reduce the failover time. Creating an Aurora replica can also help to reduce the failover time as it can take over as the primary DB instance in case of a failure. Using RDS proxy can also help to reduce the failover time as it can route the queries to the healthy DB instance, it also helps to balance the load across multiple DB instances.

NO.31 A company hosts an intranet web application on Amazon EC2 instances behind an Application Load Balancer (ALB). Currently, users authenticate to the application against an internal user database.
The company needs to authenticate users to the application by using an existing AWS Directory Service for Microsoft Active Directory directory. All users with accounts in the directory must have access to the application.
Which solution will meet these requirements?
A. Create a new app client in the directory. Create a listener rule for the ALB. Specify the authenticate-oidc action for the listener rule. Configure the listener rule with the appropriate issuer,

client ID and secret, and endpoint details for the Active Directory service. Configure the new app client with the callback URL that the ALB provides.

B. Configure an Amazon Cognito user pool. Configure the user pool with a federated identity provider (IdP) that has metadata from the directory. Create an app client. Associate the app client with the user pool. Create a listener rule for the ALB. Specify the authenticate-cognito action for the listener rule.

Configure the listener rule to use the user pool and app client.

C. Add the directory as a new 1AM identity provider (IdP). Create a new 1AM role that has an entity type of SAML 2.0 federation. Configure a role policy that allows access to the ALB. Configure the new role as the default authenticated user role for the IdP. Create a listener rule for the ALB. Specify the authenticate-oidc action for the listener rule.

D. Enable AWS 1AM Identity Center (AWS Single Sign-On). Configure the directory as an external identity provider (IdP) that uses SAML. Use the automatic provisioning method. Create a new 1AM role that has an entity type of SAML 2.0 federation. Configure a role policy that allows access to the ALB.

Attach the new role to all groups. Create a listener rule for the ALB. Specify the authenticate-cognito action for the listener rule.

Answer: A

Explanation:

The correct solution is to use the authenticate-oidc action for the ALB listener rule and configure it with the details of the AWS Directory Service for Microsoft Active Directory directory. This way, the ALB can use OpenID Connect (OIDC) to authenticate users against the directory and grant them access to the intranet web application. The app client in the directory is used to register the ALB as an OIDC client and provide the necessary credentials and endpoints. The callback URL is the URL that the ALB redirects the user to after a successful authentication. This solution does not require any additional services or roles, and it leverages the existing directory accounts for all users.

The other solutions are incorrect because they either use the wrong action for the ALB listener rule, or they involve unnecessary or incompatible services or roles. For example:

Solution B is incorrect because it uses Amazon Cognito user pool, which is a separate user directory service that does not integrate with AWS Directory Service for Microsoft Active Directory. To use this solution, the company would have to migrate or synchronize their users from the directory to the user pool, which is not required by the question. Moreover, the authenticate-cognito action for the ALB listener rule only works with Amazon Cognito user pools, not with federated identity providers (IdPs) that have metadata from the directory.

Solution C is incorrect because it uses IAM as an identity provider (IdP), which is not compatible with AWS Directory Service for Microsoft Active Directory. IAM can only be used as an IdP for web identity federation, which allows users to sign in with social media or other third-party IdPs, not with Active Directory. Moreover, the authenticate-oidc action for the ALB listener rule requires an OIDC IdP, not a SAML 2.0 federation IdP, which is what IAM provides.

Solution D is incorrect because it uses AWS IAM Identity Center (AWS Single Sign-On), which is a service that simplifies the management of SSO access to multiple AWS accounts and business applications. This service is not needed for the scenario in the question, which only involves a single intranet web application. Moreover, the authenticate-cognito action for the ALB listener rule does not work with external IdPs that use SAML, such as AWS IAM Identity Center.

References:

Authenticate users using an Application Load Balancer
What is AWS Directory Service for Microsoft Active Directory?
Using OpenID Connect for user authentication

NO.32 A company wants to migrate its website from an on-premises data center onto AWS. At the same time, it wants to migrate the website to a containerized microservice-based architecture to improve the availability and cost efficiency. The company's security policy states that privileges and network permissions must be configured according to best practice, using least privilege.
A Solutions Architect must create a containerized architecture that meets the security requirements and has deployed the application to an Amazon ECS cluster.
What steps are required after the deployment to meet the requirements? (Choose two.)
A. Create tasks using the bridge network mode.
B. Create tasks using the awsvpc network mode.
C. Apply security groups to Amazon EC2 instances, and use IAM roles for EC2 instances to access other resources.
D. Apply security groups to the tasks, and pass IAM credentials into the container at launch time to access other resources.
E. Apply security groups to the tasks, and use IAM roles for tasks to access other resources.
Answer: B E
Explanation: The awsvpc network mode provides each task with its own elastic network interface (ENI) and a primary private IP address1. By using this network mode, the solutions architect can isolate the tasks from each other and apply security groups to the tasks directly2. This way, the solutions architect can control the inbound and outbound traffic at the task level and enforce the least privilege principle3. IAM roles for tasks allow the solutions architect to assign permissions to each task separately, so that they can access other AWS resources that they need4. By using IAM roles for tasks, the solutions architect can avoid passing IAM credentials into the container at launch time, which is less secure and more prone to errors5.
References:
awsvpc network mode
Task networking with the awsvpc network mode
Security groups for your VPC
IAM roles for tasks
Best practices for managing AWS access keys

NO.33 A company manufactures smart vehicles. The company uses a custom application to collect vehicle data. The vehicles use the MQTT protocol to connect to the application.
The company processes the data in 5-minute intervals. The company then copies vehicle telematics data to on-premises storage. Custom applications analyze this data to detect anomalies.
The number of vehicles that send data grows constantly. Newer vehicles generate high volumes of data. The on-premises storage solution is not able to scale for peak traffic, which results in data loss. The company must modernize the solution and migrate the solution to AWS to resolve the scaling challenges.
Which solution will meet these requirements with the LEAST operational overhead?
A. Use AWS IOT Greengrass to send the vehicle data to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Create an Apache Kafka application to store the data in Amazon S3. Use a

pretrained model in Amazon SageMaker to detect anomalies.

B. Use AWS IOT Core to receive the vehicle data. Configure rules to route data to an Amazon Kinesis Data Firehose delivery stream that stores the data in Amazon S3. Create an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies.

C. Use AWS IOT FleetWise to collect the vehicle data. Send the data to an Amazon Kinesis data stream.

Use an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use the built-in machine learning transforms in AWS Glue to detect anomalies.

D. Use Amazon MQ for RabbitMQ to collect the vehicle data. Send the data to an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use Amazon Lookout for Metrics to detect anomalies.

Answer: B

Explanation:

Using AWS IoT Core to receive the vehicle data will enable connecting the smart vehicles to the cloud using the MQTT protocol1. AWS IoT Core is a platform that enables you to connect devices to AWS Services and other devices, secure data and interactions, process and act upon device data, and enable applications to interact with devices even when they are offline2. Configuring rules to route data to an Amazon Kinesis Data Firehose delivery stream that stores the data in Amazon S3 will enable processing and storing the vehicle data in a scalable and reliable way3. Amazon Kinesis Data Firehose is a fully managed service that delivers real-time streaming data to destinations such as Amazon S3. Creating an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies will enable analyzing the vehicle data using SQL queries or Apache Flink applications. Amazon Kinesis Data Analytics is a fully managed service that enables you to process and analyze streaming data using SQL or Java.

NO.34 A company is designing a new website that hosts static content. The website will give users the ability to upload and download large files. According to company requirements, all data must be encrypted in transit and at rest. A solutions architect is building the solution by using Amazon S3 and Amazon CloudFront.

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

A. Turn on S3 server-side encryption for the S3 bucket that the web application uses.

B. Add a policy attribute of "aws:SecureTransport": "true" for read and write operations in the S3 ACLs.

C. Create a bucket policy that denies any unencrypted operations in the S3 bucket that the web application uses.

D. Configure encryption at rest on CloudFront by using server-side encryption with AWS KMS keys (SSE-KMS).

E. Configure redirection of HTTP requests to HTTPS requests in CloudFront.

F. Use the RequireSSL option in the creation of presigned URLs for the S3 bucket that the web application uses.

Answer: A C E

Explanation:

Turning on S3 server-side encryption for the S3 bucket that the web application uses will enable encrypting the data at rest using Amazon S3 managed keys (SSE-S3)1. Creating a bucket policy that denies any unencrypted operations in the S3 bucket that the web application uses will enable

enforcing encryption for all requests to the bucket2. Configuring redirection of HTTP requests to HTTPS requests in CloudFront will enable encrypting the data in transit using SSL/TLS3.

NO.35 A company is migrating its infrastructure to the AWS Cloud. The company must comply with a variety of regulatory standards for different projects. The company needs a multi-account environment.
A solutions architect needs to prepare the baseline infrastructure. The solution must provide a consistent baseline of management and security, but it must allow flexibility for different compliance requirements within various AWS accounts. The solution also needs to integrate with the existing on-premises Active Directory Federation Services (AD FS) server.
Which solution meets these requirements with the LEAST amount of operational overhead?
A. Create an organization in AWS Organizations. Create a single SCP for least privilege access across all accounts. Create a single OU for all accounts. Configure an IAM identity provider for federation with the on-premises AD FS server. Configure a central logging account with a defined process for log generating services to send log events to the central account. Enable AWS Config in the central account with conformance packs for all accounts.
B. Create an organization in AWS Organizations. Enable AWS Control Tower on the organization. Review included controls (guardrails) for SCPs. Check AWS Config for areas that require additions. Add OUS as necessary. Connect AWS IAM Identity Center (AWS Single Sign-On) to the on-premises AD FS server.
C. Create an organization in AWS Organizations. Create SCPs for least privilege access. Create an OU structure, and use it to group AWS accounts. Connect AWS IAM Identity Center (AWS Single Sign-On) to the on-premises AD FS server. Configure a central logging account with a defined process for log generating services to send log events to the central account. Enable AWS Config in the central account with aggregators and conformance packs.
D. Create an organization in AWS Organizations. Enable AWS Control Tower on the organization. Review included controls (guardrails) for SCPs. Check AWS Config for areas that require additions. Configure an IAM identity provider for federation with the on-premises AD FS server.
Answer: B