SAP-C01 Dumps

AWS Certified Solutions Architect- Professional

https://www.certleader.com/SAP-C01-dumps.html

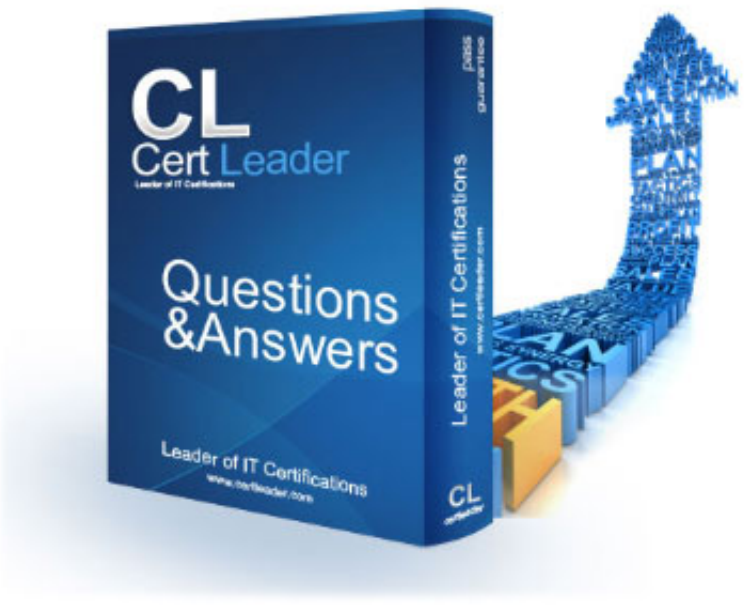

NEW QUESTION 1
A Solutions Architect is building a containerized NET Core application that will run in AWS Fargate The backend of the application requires Microsoft SQL Server with high availability All tiers of the application must be highly available The credentials used for the connection string to SQL Server should not be stored on disk within the .NET Core front-end containers.
Which strategies should the Solutions Architect use to meet these requirements'?

A. Set up SQL Server to run in Fargate with Service Auto Scaling Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to SQL Server running in Fargate Specify the ARN of the secret in AWS Secrets Manager m the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection string Set up the NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones
B. Create a Multi-AZ deployment of SQL Server on Amazon RDS Create a secret in AWS Secrets Manager for the credentials to the RDS database Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to the RDS database in Secrets Manager Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection string Set up the NET Core service m Fargate using Service Auto Scalina behind an Application Load Balancer in multiple Availability Zones.
C. Create an Auto Scaling group to run SQL Server on Amazon EC2 Create a secret in AWS Secrets Manager for the credentials to SQL Server running on EC2 Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to SQL Server on EC2 Specify the ARN of the secret in Secrets Manager In the secrets section of the Fargate task definition sothe sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection string Set up the NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availabilitv Zones.
D. Create a Multi-AZ deployment of SQL Server on Amazon RDS Create a secret in AWS Secrets Manager for the credentials to the RDS database Create non-persistent empty storage for the NET Core containers in the Fargate task definition to store the sensitive information Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to the RDS database in Secrets Manager Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be written to the non-persistent empty storage on startup for reading into the application to construct the connection.

Answer: C

NEW QUESTION 2
A group of Amazon EC2 instances have been configured as high performance computing (HPC) cluster. The instances are running in a placement group, and are able to communicate with each other at network of up to 20 Gbps.
The cluster needs to communicate with a control EC2 instance outside of the placement group. The control instance has the same instance type and AMI as the other instances, and is configured with a public IP address.
How can the Solutions Architect improve the network speeds between the control instance and the instances in the placement group?

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

Answer: D

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

NEW QUESTION 3
A company currently uses Amazon EBS and Amazon RDS for storage purposes. The company intends to use a pilot light approach for disaster recovery in a different AWS Region. The company has an RTO of 6 hours and an RPO of 24 hours.
Which solution would achieve the requirements with MINIMAL cost?

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

Answer: A

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

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

A. Create an Elasticsearch server on Amazon EC2 right-sized with 2 TB of Amazon EBS and a public AWS Elastic Beanstalk environment for the web applicatio
B. Pause the data sources, export the Elasticsearch index from on premises, and import into the EC2 Elasticsearch serve
C. Move data source feeds to the new Elasticsearch server and move users to the web application.

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

Answer: D

NEW QUESTION 5
A company had a tight deadline to migrate its on-premises environment to AWS. It moved over Microsoft SQL Servers and Microsoft Windows Servers using the virtual machine import/export service and rebuild other applications native to the cloud. The team created both Amazon EC2 databases and used Amazon RDS. Each team in the company was responsible for migrating their applications, and they have created individual accounts for isolation of resources. The company did not have much time to consider costs, but now it would like suggestions on reducing its AWS spend.
Which steps should a Solutions Architect take to reduce costs?

A. Enable AWS Business Support and review AWS Trusted Advisor's cost check
B. Create Amazon EC2 Auto Scaling groups for applications that experience fluctuating deman
C. Save AWS Simple Monthly Calculator reports in Amazon S3 for trend analysi
D. Create a master account under Organizations and have teams join for consolidating billing.
E. Enable Cost Explorer and AWS Business Support Reserve Amazon EC2 and Amazon RDS DB instance
F. Use Amazon CloudWatch and AWS Trusted Advisor for monitoring and to receive cost-savings suggestion
G. Create a master account under Organizations and have teams join for consolidated billing.
H. Create an AWS Lambda function that changes the instance size based on Amazon CloudWatch alarms.Reserve instances based on AWS Simple Monthly Calculator suggestion
I. Have an AWSWell-Architected framework review and apply recommendation
J. Create a master account under Organizations and have teams join for consolidated billing.
K. Create a budget and monitor for costs exceeding the budge
L. Create Amazon EC2 Auto Scaling groups for applications that experience fluctuating deman
M. Create an AWS Lambda function that changes instance sizes based on Amazon CloudWatch alarm
N. Have each team upload their bill to an Amazon S3 bucket for analysis of team spendin
O. Use Spot instances on nightly batch processing jobs.

Answer: D

NEW QUESTION 6
While debugging a backend application for an IoT system that supports globally distributed devices a Solutions Architect notices that stale data is occasionally being sent to user devices. Devices often share data, and stale data does not cause issues in most cases However device operations are disrupted when a device reads the stale data after an update
The global system has multiple identical application stacks deployed In different AWS Regions If a user device travels out of its home geographic region it will always connect to the geographically closest AWS Region to write or read data The same data is available in all supported AWS Regions using an Amazon DynamoDB global table
What change should be made to avoid causing disruptions in device operations'?

A. Update the backend to use strongly consistent read
B. Update the devices to always write to and read from their home AWS Region
C. Enable strong consistency globally on a DynamoDB global table Update the backend to use strongly consistent reads
D. Switch the backend data store to Amazon Aurora MySQL with cross-region replicas Update the backend to always write to the master endpoint
E. Select one AWS Region as a master and perform all writes in that AWS Region only Update the backend to use strongly consistent reads

Answer: B

NEW QUESTION 7
A company has a requirement that only allows specially hardened AMIs to be launched into public subnets in a VPC, and for the AMIs to be associated with a specific security group. Allowing non-compliant instances to launch into the public subnet could present a significant security risk if they are allowed to operate. A mapping of approved AMIs to subnets to security groups exists in an Amazon DynamoDB table in the same AWS account. The company created an AWS Lambda function that, when invoked, will terminate a given Amazon EC2 instance if the combination of AMI, subnet, and security group are not approved in the DynamoDB table.
What should the Solutions Architect do to MOST quickly mitigate the risk of compliance deviations?

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

Answer: C

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

NEW QUESTION 8
A company has a serverless application comprised of Amazon CloudFront, Amazon API Gateway, and AWS Lambda functions. The current deployment process

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

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

Answer: B

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

NEW QUESTION 9
A company CFO recently analyzed the company's AWS monthly bill and identified an opportunity to reduce the cost for AWS Elastic Beanstalk environments in use. The CFO has asked a Solutions Architect to design a highly available solution that will spin up an Elastic Beanstalk environment in the morning and terminate it at the end of the day.
The solution should be designed with minimal operational overhead and to minimize costs. It should also be able to handle the increased use of Elastic Beanstalk environments among different teams, and must provide a one-stop scheduler solution for all teams to keep the operational costs low.
What design will meet these requirements?

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

Answer: B

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

NEW QUESTION 10
A company will several AWS accounts is using AWS Organizations and service control policies (SCPs). An Administrator created the following SCP and has attached it to an organizational unit (OU) that contains AWS account 1111-1111-1111:

```
{
    "Version": "2012-10-27",
    "Statement": [
        {
            "Sid": "AllowsAllActions",
            "Effect": "Allow",
            "Action": "*",
            "Resource": "*"
        },
        {
            "Sid": "DenyCloudTrail",
            "Effect": "Deny",
            "Action": "cloudtrail:*",
            "Resource": "*"
        }
    ]
}
```

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

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

Answer: C

NEW QUESTION 10
A company is moving a business-critical, multi-tier application to AWS. The architecture consists of a desktop client application and server infrastructure. The server infrastructure resides in an on-premises data center that frequently fails to maintain the application uptime SLA of 99.95%. A Solutions Architect must re-architect the application to ensure that it can meet or exceed the SLA.
The application contains a PostgreSQL database running on a single virtual machine. The business logic and presentation layers are load balanced between multiple virtual machines. Remote users complain about slow load times while using this latency-sensitive application.
Which of the following will meet the availability requirements with little change to the application while improving user experience and minimizing costs?

A. Migrate the database to a PostgreSQL database in Amazon EC2. Host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balance
B. Allocate an Amazon WorkSpaces WorkSpace for each end user to improve the user experience.
C. Migrate the database to an Amazon RDS Aurora PostgreSQL configuratio
D. Host the application and presentation layers in an Auto Scaling configuration on Amazon EC2 instances behind an Application Load Balance
E. Use Amazon AppStream 2.0 to improve the user experience.
F. Migrate the database to an Amazon RDS PostgreSQL Multi-AZ configuratio
G. Host the application andpresentation layers in automatically scaled AWS Fargate containers behind a Network Load Balance
H. Use Amazon ElastiCache to improve the user experience.
I. Migrate the database to an Amazon Redshift cluster with at least two node
J. Combine and host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balance
K. Use Amazon CloudFront to improve the user experience.

Answer: B

NEW QUESTION 12
A company is migrating a subset of its application APIs from Amazon EC2 instances to run on a serverless infrastructure. The company has set up Amazon API Gateway, AWS Lambda, and Amazon DynamoDB for the new application. The primary responsibility of the Lambda function is to obtain data from a third-party Software as a Service (SaaS) provider. For consistency, the Lambda function is attached to the same virtual private cloud (VPC) as the original EC2 instances.
Test users report an inability to use this newly moved functionality, and the company is receiving 5xx errors from API Gateway. Monitoring reports from the SaaS provider shows that the requests never made it to its systems. The company notices that Amazon CloudWatch Logs are being generated by the Lambda functions.
When the same functionality is tested against the EC2 systems, it works as expected.
What is causing the issue?

A. Lambda is in a subnet that does not have a NAT gateway attached to it to connect to the SaaS provider.
B. The end-user application is misconfigured to continue using the endpoint backed by EC2 instances.
C. The throttle limit set on API Gateway is too low and the requests are not making their way through.
D. API Gateway does not have the necessary permissions to invoke Lambda.

Answer: A

NEW QUESTION 13
A company has an existing on-premises three-tier web application. The Linux web servers serve content from a centralized file share on a NAS server because the content is refreshed several times a day from various sources. The existing infrastructure is not optimized and the company would like to move to AWS in order to gain the ability to scale resources up and down in response to load. On-premises and AWS resources are connected using AWS Direct Connect.
How can the company migrate the web infrastructure to AWS without delaying the content refresh process?

A. Create a cluster of web server Amazon EC2 instances behind a Classic Load Balancer on AW
B. Share an Amazon EBS volume among all instances for the conten
C. Schedule a periodic synchronization of this volume and the NAS server.
D. Create an on-premises file gateway using AWS Storage Gateway to replace the NAS server and replicate content to AW
E. On the AWS side, mount the same Storage Gateway bucket to each web server Amazon EC2 instance to serve the content.
F. Expose an Amazon EFS share to on-premises users to serve as the NAS serv
G. Mount the same EFS share to the web server Amazon EC2 instances to serve the content.
H. Create web server Amazon EC2 instances on AWS in an Auto Scaling grou
I. Configure a nightly process where the web server instances are updated from the NAS server.

Answer: C

Explanation:
File gateway is limited by performance its gateway instance, whether EC2 or On-premises, Cache will get filled up fast if not properly configured, For large number of EC2 instances EFS scales better. So, bottom line is File Storage gateway is for legacy applications and you have to add cost of large gateway instances before comparing it to same quantity of EFS storage. https://www.reddit.com/r/aws/comments/82pyop/storage_gateway_vs_efs/
https://docs.aws.amazon.com/efs/latest/ug/efs-onpremises.html

NEW QUESTION 18
A Development team is deploying new APIs as serverless applications within a company. The team is currently using the AWS Management Console to provision Amazon API Gateway, AWS Lambda, and Amazon DynamoDB resources. A Solutions Architect has been tasked with automating the future deployments of these serverless APIs.
How can this be accomplished?

A. Use AWS CloudFormation with a Lambda-backed custom resource to provision API Gatewa

B. Use the AWS::DynamoDB::Table and AWS::Lambda::Function resources to create the Amazon DynamoDB table and Lambda function
C. Write a script to automate the deployment of the CloudFormation template.
D. Use the AWS Serverless Application Model to define the resource
E. Upload a YAML template and application files to the code repositor
F. Use AWS CodePipeline to connect to the code repository and to create an action to build using AWS CodeBuil
G. Use the AWS CloudFormation deployment provider in CodePipeline to deploy the solution.
H. Use AWS CloudFormation to define the serverless applicatio
I. Implement versioning on the Lambda functions and create aliases to point to the version
J. When deploying, configure weights to implement shifting traffic to the newest version, and gradually update the weights as traffic moves over.
K. Commit the application code to the AWS CodeCommit code repositor
L. Use AWS CodePipeline and connect to the CodeCommit code repositor
M. Use AWS CodeBuild to build and deploy the Lambda functions using AWS CodeDeplo
N. Specify the deployment preference type in CodeDeploy to gradually shift traffic over to the new version.

Answer: B

Explanation:
https://aws-quickstart.s3.amazonaws.com/quickstart-trek10-serverless-enterprise-cicd/doc/serverless-cicd-for-th
https://aws.amazon.com/quickstart/architecture/serverless-cicd-for-enterprise/

NEW QUESTION 22
A Solutions Architect is designing a network solution for a company that has applications running in a data center in Northern Virginia. The applications in the company's data center require predictable performance to applications running in a virtual private cloud (VPC) located in us-east-1, and a secondary VPC in us-west-2 within the same account. The company data center is collocated in an AWS Direct Connect facility that serves the us-est-1 region. The company has already ordered an AWS Direct Connect connection and a cross-connect has been established.
Which solution will meet the requirements at the LOWEST cost?

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

Answer: A

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

NEW QUESTION 26
A large multinational company runs a timesheet application on AWS that is used by staff across the world. The application runs on Amazon EC2 instances in an Auto Scaling group behind an Elastic Load Balancing (ELB) load balancer, and stores in an Amazon RDS MySQL Multi-AZ database instance.
The CFO is concerned about the impact on the business if the application is not available. The application must not be down for more than two hours, but the solution must be as cost-effective as possible.
How should the Solutions Architect meet the CFO's requirements while minimizing data loss?

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

Answer: D

NEW QUESTION 28
A company is using AWS for production and development workloads. Each business unit has its own AWS account for production, and a separate AWS account to develop and deploy its applications. The Information Security department has introduced new security policies that limit access for terminating certain Amazon ECs instances in all accounts to a small group of individuals from the Security team.
How can the Solutions Architect meet these requirements?

A. Create a new IAM policy that allows access to those EC2 instances only for the Security tea
B. Apply this policy to the AWS Organizations master account.
C. Create a new tag-based IAM policy that allows access to these EC2 instances only for the Security team.Tag the instances appropriately, and apply this policy

in each account.
D. Create an organizational unit under AWS Organization
E. Move all the accounts into this organizational unit and use SCP to apply a whitelist policy to allow access to these EC2 instances for the Security teamonly.
F. Set up SAML federation for all accounts in AW
G. Configure SAML so that it checks for the service API call before authenticating the use
H. Block SAML from authenticating API calls if anyone other than the Security team accesses these instances.

Answer: B

NEW QUESTION 33
A company runs a dynamic mission-critical web application that has an SLA of 99.99%. Global application users access the application 24/7. The application is currently hosted on premises and routinely fails to meet its SLA, especially when millions of users access the application concurrently. Remote users complain of latency.
How should this application be redesigned to be scalable and allow for automatic failover at the lowest cost?

A. Use Amazon Route 53 failover routing with geolocation-based routin
B. Host the website on automatically scaled Amazon EC2 instances behind an Application Load Balancer with an additional Application Load Balancer and EC2 instances for the application layer in each regio
C. Use a Multi-AZ deployment with MySQL as the data layer.
D. Use Amazon Route 53 round robin routing to distribute the load evenly to several regions with health check
E. Host the website on automatically scaled Amazon ECS with AWS Fargate technology containers behind a Network Load Balancer, with an additional Network Load Balancer and Fargate containers for the application layer in each regio
F. Use Amazon Aurora replicas for the data layer.
G. Use Amazon Route 53 latency-based routing to route to the nearest region with health check
H. Host the website in Amazon S3 in each region and use Amazon API Gateway with AWS Lambda for the application laye
I. Use Amazon DynamoDB global tables as the data layer with Amazon DynamoDB Accelerator (DAX) for caching.
J. Use Amazon Route 53 geolocation-based routin
K. Host the website on automatically scaled AWS Fargate containers behind a Network Load Balancer with an additional Network Load Balancer and Fargate containers for the application layer in each regio
L. Use Amazon Aurora Multi-Master for Aurora MySQL as the data layer.

Answer: C

Explanation:
https://aws.amazon.com/getting-started/projects/build-serverless-web-app-lambda-apigateway-s3-dynamodb-co

NEW QUESTION 38
A company has an application that uses Amazon EC2 instances in an Auto Scaling group. The Quality Assurance (QA) department needs to launch a large number of short-lived environments to test the application. The application environments are currently launched by the Manager of the department using an AWS CloudFormation template. To launch the stack, the Manager uses a role with permission to use CloudFormation, EC2 and Auto Scaling APIs. The Manager wants to allow testers to launch their own environments, but does not want to grant broad permission to each user. Which set up would achieve these goals?

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

Answer: B

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

NEW QUESTION 40
An online e-commerce business is running a workload on AWS. The application architecture includes a web tier, an application tier for business logic, and a database tier for user and transactional data management. The database server has a 100 GB memory requirement. The business requires cost-efficient disaster recovery for the application with an RTO of 5 minutes and an RPO of 1 hour. The business also has a regulatory for out-of region disaster recovery with a minimum distance between the primary and alternate sites of 250 miles.
Which of the following options can the Solutions Architect design to create a comprehensive solution for this customer that meets the disaster recovery requirements?

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

I. Activate the primary database in one region only and the standby database in the other regio
J. Use Amazon Route 53 to automatically switch traffic from one region to another using health check routing policies.

Answer: C

NEW QUESTION 42
A Solutions Architect has created an AWS CloudFormation template for a three-tier application that contains an Auto Scaling group of Amazon EC2 instances running a custom AMI.
The Solutions Architect wants to ensure that future updates to the custom AMI can be deployed to a running stack by first updating the template to refer to the new AMI, and then invoking UpdateStack to replace the EC2 instances with instances launched from the new AMI.
How can updates to the AMI be deployed to meet these requirements?

A. Create a change set for a new version of the template, view the changes to the running EC2 instances to ensure that the AMI is correctly updated, and then execute the change set.
B. Edit the AWS::AutoScaling::LaunchConfiguration resource in the template, changing its DeletionPolicy to Replace.
C. Edit the AWS::AutoScaling:: AutoScalingGroup resource in the template, inserting an UpdatePolicy attribute.
D. Create a new stack from the updated templat
E. Once it is successfully deployed, modify the DNS records to point to the new stack and delete the old stack.

Answer: C

Explanation:
References:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-as-launchconfig.html

NEW QUESTION 43
A company deployed a three-tier web application in two regions: us-east-1 and eu-west-1. The application must be active in both regions at the same time. The database tier of the application uses a single Amazon RDS Aurora database globally, with a master in us-east-1 and a read replica in eu-west-1. Both regions are connected by a VPN.
The company wants to ensure that the application remains available even in the event of a region-level failure of all of the application's components. It is acceptable for the application to be in read-only mode for up to 1 hour. The company plans to configure two Amazon Route 53 record sets, one for each of the regions.
How should the company complete the configuration to meet its requirements while providing the lowest latency for the application end-users? (Choose two.)

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

Answer: CE

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

NEW QUESTION 48
A media company has a 30-TB repository of digital news videos. These videos are stored on tape in an
on-premises tape library and referenced by a Media Asset Management (MAM) system. The company wants to enrich the metadata for these videos in an automated fashion and put them into a searchable catalog by using a MAM feature. The company must be able to search based on information in the video, such as objects, scenery items, or people's faces. A catalog is available that contains faces of people who have appeared in the videos that include an image of each person. The company would like to migrate these videos to AWS.
The company has a high-speed AWS Direct Connect connection with AWS and would like to move the MAM solution video content directly from its current file system.
How can these requirements be met by using the LEAST amount of ongoing management overhead and causing MINIMAL disruption to the existing system?

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

Answer: C

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

NEW QUESTION 49
A group of research institutions and hospitals are in a partnership to study 2 PBs of genomic data. The institute that owns the data stores it in an Amazon S3 bucket and updates it regularly. The institute would like to give all of the organizations in the partnership read access to the data. All members of the partnership are extremely cost-conscious, and the institute that owns the account with the S3 bucket is concerned about covering the costs for requests and data transfers from Amazon S3.
Which solution allows for secure datasharing without causing the institute that owns the bucket to assume all the costs for S3 requests and data transfers?

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

Answer: B

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

NEW QUESTION 53
A company wants to allow its Marketing team to perform SQL queries on customer records to identify market segments. The data is spread across hundreds of files. The records must be encrypted in transit and at rest. The Team Manager must have the ability to manage users and groups, but no team members should have access to services or resources not required for the SQL queries. Additionally, Administrators need to audit the queries made and receive notifications when a query violates rules defined by the Security team.
AWS Organizations has been used to create a new account and an AWS IAM user with administrator permissions for the Team Manager.
Which design meets these requirements?

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

Answer: D

NEW QUESTION 56
The Solutions Architect manages a serverless application that consists of multiple API gateways, AWS Lambda functions, Amazon S3 buckets, and Amazon DynamoDB tables. Customers say that a few application components slow while loading dynamic images, and some are timing out with the "504 Gateway Timeout" error. While troubleshooting the scenario, the Solutions Architect confirms that DynamoDB monitoring metrics are at acceptable levels.
Which of the following steps would be optimal for debugging these application issues? (Choose two.)

A. Parse HTTP logs in Amazon API Gateway for HTTP errors to determine the root cause of the errors.
B. Parse Amazon CloudWatch Logs to determine processing times for requested images at specified intervals.
C. Parse VPC Flow Logs to determine if there is packet loss between the Lambda function and S3.
D. Parse AWS X-Ray traces and analyze HTTP methods to determine the root cause of the HTTP errors.
E. Parse S3 access logs to determine if objects being accessed are from specific IP addresses to narrow the scope to geographic latency issues.

Answer: BD

Explanation:
Firstly "A 504 Gateway Timeout Error means your web server didn't receive a timely response from another server upstream when it attempted to load one of your web pages. Put simply, your web servers aren't communicating with each other fast enough". This specific issue is addressed in the AWS article "Tracing, Logging and Monitoring an API Gateway API". https://docs.amazonaws.cn/en_us/apigateway/latest/developerguide/monitoring_overview.html

NEW QUESTION 60
A company has been using a third-party provider for its content delivery network and recently decided to switch to Amazon CloudFront the Development team

wants to maximize performance for the global user base. The company uses a content management system (CMS) that serves both static and dynamic content. The CMS is both md an Application Load Balancer (ALB) which is set as the default origin for the distribution. Static assets are served from an Amazon S3 bucket. The Origin Access Identity (OAI) was created property d the S3 bucket policy has been updated to allow the GetObject action from the OAI, but static assets are receiving a 404 error
Which combination of steps should the Solutions Architect take to fix the error? (Select TWO.)

A. Add another origin to the CloudFront distribution for the static assets
B. Add a path based rule to the ALB to forward requests for the static assets
C. Add an RTMP distribution to allow caching of both static and dynamic content
D. Add a behavior to the CloudFront distribution for the path pattern and the origin of the static assets
E. Add a host header condition to the ALB listener and forward the header from CloudFront to add traffic to the allow list

Answer: AD

NEW QUESTION 62
A company is using AWS to run an internet-facing production application written in Node.js. The Development team is responsible for pushing new versions of their software directly to production. The application software is updated multiple times a day. The team needs guidance from a Solutions Architect to help them deploy the software to the production fleet quickly and with the least amount of disruption to the service.
Which option meets these requirements?

A. Prepackage the software into an AMI and then use Auto Scaling to deploy the production flee
B. For software changes, update the AMI and allow Auto Scaling to automatically push the new AMI to production.
C. Use AWS CodeDeploy to push the prepackaged AMI to productio
D. For software changes, reconfigure CodeDeploy with new AMI identification to push the new AMI to the production fleet.
E. Use AWS Elastic Beanstalk to host the production applicatio
F. For software changes, upload the new application version to Elastic Beanstalk to push this to the production fleet using a blue/green deployment method.
G. Deploy the base AMI through Auto Scaling and bootstrap the software using user dat
H. For software changes, SSH to each of the instances and replace the software with the new version.

Answer: C

NEW QUESTION 63
A company has a legacy application running on servers on premises. To increase the application's reliability, the company wants to gain actionable insights using application logs. A Solutions Architect has been given following requirements for the solution:

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

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

Answer: A

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

NEW QUESTION 68
An enterprise company is using a multi-account AWS strategy There are separate accounts tor development staging and production workloads To control costs and improve governance the following requirements have been defined:
• The company must be able to calculate the AWS costs tor each project
• The company must be able to calculate the AWS costs tor each environment development staging and production
• Commonly deployed IT services must be centrally managed
• Business units can deploy pre-approved IT services only
• Usage of AWS resources in the development account must be limited
Which combination of actions should be taken to meet these requirements? (Select THREE)

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

Answer: CEF

NEW QUESTION 71
A company has several teams, and each team has their own Amazon RDS database that totals 100 TB The company is building a data query platform for Business Intelligence Analysts to generate a weekly business report The new system must run ad-hoc SQL queries
What is the MOST cost-effective solution?

A. Create a new Amazon Redshift cluster Create an AWS Glue ETL job to copy data from the RDS databases to the Amazon Redshift cluster Use Amazon

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

Answer: C

NEW QUESTION 72
A company is migrating to the cloud. It wants to evaluate the configurations of virtual machines in its existing data center environment to ensure that it can size new Amazon EC2 instances accurately. The company wants to collect metrics, such as CPU, memory, and disk utilization, and it needs an inventory of what processes are running on each instance. The company would also like to monitor network connections to map communications between servers.
Which would enable the collection of this data MOST cost effectively?

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

Answer: A

NEW QUESTION 76
A Solutions Architect must create a cost-effective backup solution for a company's 500MB source code repository of proprietary and sensitive applications. The repository runs on Linux and backs up daily to tape. Tape backups are stored for 1 year.
The current solutions are not meeting the company's needs because it is a manual process that is prone to error, expensive to maintain, and does not meet the need for a Recovery Point Objective (RPO) of 1 hour or Recovery Time Objective (RTO) of 2 hours. The new disaster recovery requirement is for backups to be stored offsite and to be able to restore a single file if needed.
Which solution meets the customer's needs for RTO, RPO, and disaster recovery with the LEAST effort and expense?

A. Replace local tapes with an AWS Storage Gateway virtual tape library to integrate with current backup softwar
B. Run backups nightly and store the virtual tapes on Amazon S3 standard storage inUS-EAST-1. Use cross-region replication to create a second copy in US-WEST-2. Use Amazon S3 lifecycle policies to perform automatic migration to Amazon Glacier and deletion of expired backups after 1 year?
C. Configure the local source code repository to synchronize files to an AWS Storage Gateway file Amazon gateway to store backup copies in an Amazon S3 Standard bucke
D. Enable versioning on the Amazon S3 bucke
E. Create Amazon S3 lifecycle policies to automatically migrate old versions of objects to Amazon S3 Standard 0 Infrequent Access, then Amazon Glacier, then delete backups after 1 year.
F. Replace the local source code repository storage with a Storage Gateway stored volum
G. Change the default snapshot frequency to 1 hou
H. Use Amazon S3 lifecycle policies to archive snapshots to Amazon Glacier and remove old snapshots after 1 yea
I. Use cross-region replication to create a copy of the snapshots in US-WEST-2.
J. Replace the local source code repository storage with a Storage Gateway cached volum
K. Create a snapshot schedule to take hourly snapshot
L. Use an Amazon CloudWatch Events schedule expression rule to run on hourly AWS Lambda task to copy snapshots from US-EAST -1 to US-WEST-2.

Answer: B

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

NEW QUESTION 78
A company manages more than 200 separate internet-facing web applications. All of the applications are deployed to AWS in a single AWS Region The fully qualified domain names (FQDNs) of all of the applications are made available through HTTPS using Application Load Balancers (ALBs). The ALBs are configured to use public SSL/TLS certificates.
A Solutions Architect needs to migrate the web applications to a multi-region architecture. All HTTPS services should continue to work without interruption.
Which approach meets these requirements?

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

Answer: D

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

NEW QUESTION 80
A company has more than 100 AWS accounts, with one VPC per account, that need outbound HTTPS connectivity to the internet. The current design contains one

NAT gateway per Availability Zone (AZ) in each VPC. To reduce costs and obtain information about outbound traffic, management has asked for a new architecture for internet access.
Which solution will meet the current needs, and continue to grow as new accounts are provisioned, while reducing costs?

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

Answer: D

Explanation:
user proxy fleet over PrivateLink. As explained in this AWS website:
https://aws.amazon.com/blogs/networking-and-content-delivery/how-to-use-aws-privatelink-to-secure-and-scale

NEW QUESTION 82
A company's data center is connected to the AWS Cloud over a minimally used 10-Gbps AWS Direct Connect connection with a private virtual interface to its virtual private cloud {VPC) The company internet connection is 200 Mbps and the company has a 150-TB dataset that is created each Friday The data must be transferred and available in Amazon S3 on Monday morning
Which is the LEAST expensive way to meet the requirements while allowing for data transfer growth?

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

Answer: D

NEW QUESTION 87
A company operating a website on AWS requires high levels of scalability, availability and performance. The company is running a Ruby on Rails application on Amazon EC2. It has a data tier on MySQL 5.6 on Amazon EC2 using 16 TB of Amazon EBS storage. Amazon CloudFront is used to cache application content. The Operations team is reporting continuous and unexpected growth of EBS volumes assigned to the MySQL database. The Solutions Architect has been asked to design a highly scalable, highly available, and high-performing solution.
Which solution is the MOST cost-effective at scale?

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

Answer: C

NEW QUESTION 91
A Solutions Architect must establish a patching plan for a large mixed fleet of Windows and Linux servers. The patching plan must be implemented securely, be audit ready, and comply with the company's business requirements.
Which option will meet these requirements with MINIMAL effort?

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

Answer: B

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

NEW QUESTION 94
An e-commerce company is revamping its IT infrastructure and is planning to use AWS services. The company's CIO has asked a Solutions Architect to design a simple, highly available, and loosely coupled order processing application. The application is responsible for receiving and processing orders before storing them in an Amazon DynamoDB table. The application has a sporadic traffic pattern and should be able to scale during marketing campaigns to process the orders with minimal delays.
Which of the following is the MOST reliable approach to meet the requirements?

A. Receive the orders in an Amazon EC2-hosted database and use EC2 instances to process them.
B. Receive the orders in an Amazon SQS queue and trigger an AWS Lambda function to process them.
C. Receive the orders using the AWS Step Functions program and trigger an Amazon ECS container to process them.
D. Receive the orders in Amazon Kinesis Data Streams and use Amazon EC2 instances to process them.

Answer: B

NEW QUESTION 97
A company has a large on-premises Apache Hadoop cluster with a 20 PB HDFS database. The cluster is growing every quarter by roughly 200 instances and 1 PB. The company's goals are to enable resiliency for its Hadoop data, limit the impact of losing cluster nodes, and significantly reduce costs. The current cluster runs 24/7 and supports a variety of analysis workloads, including interactive queries and batch processing.
Which solution would meet these requirements with the LEAST expense and down time?

A. Use AWS Snowmobile to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workload based on historical data from theon-premises cluste
B. Store the data on EMRF
C. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metric
D. Create job-specific, optimized clusters for batch workloads that are similarly optimized.
E. Use AWS Snowmobile to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster of similar size and configuration to the current cluste
F. Store the data on EMRF
G. Minimize costs by using Reserved Instance
H. As the workload grows each quarter, purchase additional Reserved Instances and add to the cluster.
I. Use AWS Snowball to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workloads based on historical data from theon-premises cluste
J. Store the on EMRF
K. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metric
L. Create job-specific, optimized clusters for batch workloads that are similarly optimized.
M. Use AWS Direct Connect to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workload based on historical data from theon-premises cluste
N. Store the data on EMRF
O. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metric
P. Create job-specific, optimized clusters for batch workloads that are similarly optimized.

Answer: A

Explanation:
Q: How should I choose between Snowmobile and Snowball?
To migrate large datasets of 10PB or more in a single location, you should use Snowmobile. For datasets less than 10PB or distributed in multiple locations, you should use Snowball. In addition, you should evaluate the amount of available bandwidth in your network backbone. If you have a high speed backbone with hundreds of Gb/s of spare throughput, then you can use Snowmobile to migrate the large datasets all at once. If you have limited bandwidth on your backbone, you should consider using multiple Snowballs to migrate the data incrementally.

NEW QUESTION 98
A Solutions Architect must migrate an existing on-premises web application with 70 TB of static files supporting a public open-data initiative. The architect wants to upgrade to the latest version of the host operating system as part of the migration effort.
Which is the FASTEST and MOST cost-effective way to perform the migration?

A. Run a physical-to-virtual conversion on the application serve
B. Transfer the server image over the internet, and transfer the static data to Amazon S3.
C. Run a physical-to-virtual conversion on the application serve
D. Transfer the server image over AWS Direct Connect, and transfer the static data to Amazon S3.
E. Re-platform the server to Amazon EC2, and use AWS Snowball to transfer the static data to Amazon S3.
F. Re-platform the server by using the AWS Server Migration Service to move the code and data to a new Amazon EC2 instance.

Answer: C

NEW QUESTION 102
A company wants to replace its call system with a solution built using AWS managed services. The company call center would like the solution to receive calls, create contact flows, and scale to handle growth projections. The call center would also like the solution to use deep learning capabilities to recognize the intent of the callers and handle basic tasks, reducing the need to speak an agent. The solution should also be able to query business applications and provide relevant information back to calls as requested.
Which services should the Solution Architect use to build this solution? (Choose three.)

A. Amazon Rekognition to identity who is calling.
B. Amazon Connect to create a cloud-based contact center.
C. Amazon Alexa for Business to build conversational interface.
D. AWS Lambda to integrate with internal systems.
E. Amazon Lex to recognize the intent of the caller.

F. Amazon SQS to add incoming callers to a queue.

Answer: BDE

NEW QUESTION 107
A company's main intranet page has experienced degraded response times as its user base has increased although there are no reports of users seeing error pages. The application uses Amazon DynamoDB in read-only mode.
Amazon DynamoDB latency metrics for successful requests have been in a steady state even during times when users have reported degradation The Development team has correlated the issue to ProvisionedThrough put Exceeded exceptions in the application logs when doing Scan and read operations The team also identified an access pattern of steady spikes of read activity on a distributed set of individual data items
The Chief Technology Officer wants to improve the user experience
Which solutions will meet these requirements with the LEAST amount of changes to the application? (Select TWO)

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

Answer: AE

NEW QUESTION 109
A company plans to move regulated and security-sensitive businesses to AWS. The Security team is developing a framework to validate the adoption of AWS best practice and industry-recognized compliance standards. The AWS Management Console is the preferred method for teams to provision resources.
Which strategies should a Solutions Architect use to meet the business requirements and continuously assess, audit, and monitor the configurations of AWS resources? (Choose two.)

A. Use AWS Config rules to periodically audit changes to AWS resources and monitor the compliance of the configuratio
B. Develop AWS Config custom rules using AWS Lambda to establish a test-driven development approach, and further automate the evaluation of configuration changes against the required controls.
C. Use Amazon CloudWatch Logs agent to collect all the AWS SDK log
D. Search the log data using a pre-defined set of filter patterns that machines mutating API call
E. Send notifications using Amazon CloudWatch alarms when unintended changes are performe
F. Archive log data by using a batch exportto Amazon S3 and then Amazon Glacier for a long-term retention and auditability.
G. Use AWS CloudTrail events to assess management activities of all AWS account
H. Ensure that CloudTrail is enabled in all accounts and available AWS service
I. Enable trails, encrypt CloudTrail event log files with an AWS KMS key, and monitor recorded activities with CloudWatch Logs.
J. Use the Amazon CloudWatch Events near-real-time capabilities to monitor system events patterns, and trigger AWS Lambda functions to automatically revert non-authorized changes in AWS resource
K. Also, target Amazon SNS topics to enable notifications and improve the response time of incident responses.
L. Use CloudTrail integration with Amazon SNS to automatically notify unauthorized API activities.Ensure that CloudTrail is enabled in all accounts and available AWS service
M. Evaluate the usage of Lambda functions to automatically revert non-authorized changes in AWS resources.

Answer: AC

Explanation:
https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudwatch-alarms-for-cloudtrail.html
https://docs.aws.amazon.com/en_pv/awscloudtrail/latest/userguide/best-practices-security.html

NEW QUESTION 112
The CISO of a large enterprise with multiple IT departments, each with its own AWS account, wants one central place where AWS permissions for users can be managed and users authentication credentials can be synchronized with the company’s existing on-premises solution.
Which solution will meet the CISO’s requirements?

A. Define AWS IAM roles based on the functional responsibilities of the users in a central accoun
B. Create a SAML-based identity management provide
C. Map users in the on-premises groups to IAM role
D. Establish trust relationships between the other accounts and the central account.
E. Deploy a common set of AWS IAM users, groups, roles, and policies in all of the AWS accounts using AWS Organization
F. Implement federation between the on-premises identity provider and the AWS accounts.
G. Use AWS Organizations in a centralized account to define service control policies (SCPs). Create a SAML-based identity management provider in each account and map users in the on-premises groups to AWS IAM roles.
H. Perform a thorough analysis of the user base and create AWS IAM users accounts that have the necessary permission
I. Set up a process to provision and de provision accounts based on data in the on-premises solution.

Answer: A

Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 114
A company needs to run a software package that has a license that must be run on the same physical host for the duration of its use. The software package is only going to be used for 90 days. The company requires patching and restarting of all instances every 30 days.
How can these requirements be met using AWS?

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

Answer: B

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

NEW QUESTION 119
A Solutions Architect is working with a company that operates a standard three-tier web application in AWS. The web and application tiers run on Amazon EC2 and the database tier runs on Amazon RDS. The company is redesigning the web and application tiers to use Amazon API Gateway and AWS Lambda, and the company intends to deploy the new application within 6 months. The IT Manager has asked the Solutions Architect to reduce costs in the interim.
Which solution will be MOST cost effective while maintaining reliability?

A. Use Spot Instances for the web tier, On-Demand Instances for the application tier, and Reserved Instances for the database tier.
B. Use On-Demand Instances for the web and application tiers, and Reserved Instances for the database tier.
C. Use Spot Instances for the web and application tiers, and Reserved Instances for the database tier.
D. Use Reserved Instances for the web, application, and database tiers.

Answer: B

NEW QUESTION 120
A retail company is running an application that stores invoice files in Amazon S3 bucket and metadata about the files in an Amazon. The S3 bucket and DynamoDB table are in us-east-1. The company wants to protect itself from data corruption and loss of connectivity to either Region.
Which option meets these requirements?

A. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable continuous backup on the DynamoDB table in us-east-1. Enable versioning on the S3 bucket.
B. Create an AWS Lambda function triggered by Amazon CloudWatch Events to make regular backups of the DynamoDB tabl
C. Set up S3 cross-region replication from us-east-1 to eu-west-1. Set up MFA deleteon the S3 bucket in us-east-1.
D. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable versioning on the S3 bucke
E. Implement strict ACLs on the S3 bucket.
F. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable continuous backup on the DynamoDB table in us-east-1. Set up S3 cross-region replication from us-east-1 toeu-west-1.

Answer: A

Explanation:
https://aws.amazon.com/blogs/aws/new-cross-region-replication-for-amazon-s3/

NEW QUESTION 122
An online retailer needs to regularly process large product catalogs, which are handled in batches. These are sent out to be processed by people using the Amazon Mechanical Turk service, but the retailer has asked its Solutions Architect to design a workflow orchestration system that allows it to handle multiple concurrent Mechanical Turk operations, deal with the result assessment process, and reprocess failures.
Which of the following options gives the retailer the ability to interrogate the state of every workflow with the LEAST amount of implementation effort?

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

Answer: C

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

NEW QUESTION 125
A company has implemented AWS Organizations. It has recently set up a number of new accounts and wants to deny access to a specific set of AWS services in these new accounts.
How can this be controlled MOST efficiently?

A. Create an IAM policy in each account that denies access to the service
B. Associate the policy with an IAM group, and add all IAM users to the group.
C. Create a service control policy that denies access to the service

D. Add all of the new accounts to a single organizations unit (OU), and apply the policy to that OU.
E. Create an IAM policy in each account that denies access to the servic
F. Associate the policy with an IAM role, and instruct users to log in using their corporate credentials and assume the IAM role.
G. Create a service control policy that denies access to the services, and apply the policy to the root of the organization.

Answer: B

NEW QUESTION 130
A company has an application behind a load balancer with enough Amazon EC2 instances to satisfy peak demand. Scripts and third-party deployment solutions are used to configure EC2 instances when demand increases or an instance fails. The team must periodically evaluate the utilization of the instance types to ensure that the correct sizes are deployed.
How can this workload be optimized to meet these requirements?

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

Answer: D

NEW QUESTION 133
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: BE

Explanation:
https://aws.amazon.com/about-aws/whats-new/2017/11/amazon-ecs-introduces-awsvpc-networking-mode-for-c
https://amazonaws-china.com/blogs/compute/introducing-cloud-native-networking-for-ecs-containers/
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-iam-roles.html

NEW QUESTION 135
A company currently uses a single 1 Gbps AWS Direct Connect connection to establish connectivity between an AWS Region and its data center. The company has five Amazon VPCs, all of which are connected to the data center using the same Direct Connect connection. The Network team is worried about the single point of failure and is interested in improving the redundancy of the connections to AWS while keeping costs to a minimum.
Which solution would improve the redundancy of the connection to AWS while meeting the cost requirements?

A. Provision another 1 Gbps Direct Connect connection and create new VIFs to each of the VPCs.Configure the VIFs in a load balancing fashion using BGP.
B. Set up VPN tunnels from the data center to each VP
C. Terminate each VPN tunnel at the virtual private gateway (VGW) of the respective VPC and set up BGP for route management.
D. Set up a new point-to-point Multiprotocol Label Switching (MPLS) connection to the AWS Region that's being use
E. Configure BGP to use this new circuit as passive, so that no traffic flows through this unless the AWS Direct Connect fails.
F. Create a public VIF on the Direct Connect connection and set up a VPN tunnel which will terminate on the virtual private gateway (VGW) of the respective VPC using the public VI
G. Use BGP to handle the failover to the VPN connection.

Answer: B

NEW QUESTION 137
A Solutions Architect needs to design a highly available application that will allow authenticated users to stay connected to the application even when there are underlying failures
Which solution will meet these requirements?

A. Deploy the application on Amazon EC2 instances Use Amazon Route 53 to forward requests to the EC2 Instances Use Amazon DynamoDB to save the authenticated connection details
B. Deploy the application on Amazon EC2 instances in an Auto Scaling group Use an internet-facing Application Load Balancer to handle requests Use Amazon DynamoDB to save the authenticated connection details
C. Deploy the application on Amazon EC2 instances in an Auto Scaling group Use an internet-facing Application Load Balancer on the front end Use EC2 instances to save the authenticated connectiondetails
D. Deploy the application on Amazon EC2 instances in an Auto Scaling group Use an internet-facing Application Load Balancer on the front end Use EC2 instances hosting a MySQL database to save the authenticated connection details

Answer: B

NEW QUESTION 141
A media storage application uploads user photos to Amazon S3 for processing. End users are reporting that some uploaded photos are not being processed properly. The Application Developers trace the logs and find that AWS Lambda is experiencing execution issues when thousands of users are on the system simultaneously. Issues are caused by:

- Limits around concurrent executions.
- The performance of Amazon DynamoDB when saving data.

Which actions can be taken to increase the performance and reliability of the application? (Choose two.)

A. Evaluate and adjust the read capacity units (RCUs) for the DynamoDB tables.
B. Evaluate and adjust the write capacity units (WCUs) for the DynamoDB tables.
C. Add an Amazon ElastiCache layer to increase the performance of Lambda functions.
D. Configure a dead letter queue that will reprocess failed or timed-out Lambda functions.
E. Use S3 Transfer Acceleration to provide lower-latency access to end users.

Answer: BD

Explanation:
B:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.h
D: https://aws.amazon.com/blogs/compute/robust-serverless-application-design-with-aws-lambda-dlq/c

NEW QUESTION 146
A company has deployed an application to multiple environments in AWS, including production and testing. The company has separate accounts for production and testing, and users are allowed to create additional application users for team members or services, as needed. The Security team has asked the Operations team for better isolation between production and testing with centralized controls on security credentials and improved management of permissions between environments.
Which of the following options would MOST securely accomplish this goal?

A. Create a new AWS account to hold user and service accounts, such as an identity accoun
B. Create users and groups in the identity accoun
C. Create roles with appropriate permissions in the production and testing account
D. Add the identity account to the trust policies for the roles.
E. Modify permissions in the production and testing accounts to limit creating new IAM users to members of the Operations tea
F. Set a strong IAM password policy on each accoun
G. Create new IAM users and groups in each account to limit developer access to just the services required to complete their job function.
H. Create a script that runs on each account that checks user accounts for adherence to a security policy.Disable any user or service accounts that do not comply.
I. Create all user accounts in the production accoun
J. Create roles for access in the production account and testing account
K. Grant cross-account access from the production account to the testing account.

Answer: A

Explanation:
https://aws.amazon.com/blogs/security/how-to-centralize-and-automate-iam-policy-creation-in-sandbox-develop

NEW QUESTION 148
A company is building an AWS landing zone and has asked a Solutions Architect to design a multi-account access strategy that will allow hundreds of users to use corporate credentials to access the AWS Console. The company is running a Microsoft Active Directory and users will use an AWS Direct Connect connection to connect to AWS. The company also wants to be able to federate to third-party services and providers, including custom applications.
Which solution meets the requirements by using the LEAST amount of management overhead?

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

Answer: D

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

NEW QUESTION 149
An auction website enables users to bid on collectible items. The auction rules require that each bid is processed only once and in the order it was received. The current implementation is based on a fleet of Amazon EC2 web servers that write bid records into Amazon Kinesis Data Streams. A single t2.large instance has a cron job that runs the bid processor, which reads incoming bids from Kinesis Data Streams and processes each bid. The auction site is growing in popularity, but users are complaining that some bids are not registering.
Troubleshooting indicates that the bid processor is too slow during peak demand hours, sometimes crashes while processing, and occasionally loses track of

which records is being processed.
What changes should make the bid processing more reliable?

A. Refactor the web application to use the Amazon Kinesis Producer Library (KPL) when posting bids to Kinesis Data Stream
B. Refactor the bid processor to flag each record in Kinesis Data Streams as being unread, processing, and processe
C. At the start of each bid processing run, scan Kinesis Data Streams for unprocessed records.
D. Refactor the web application to post each incoming bid to an Amazon SNS topic in place of Kinesis Data Stream
E. Configure the SNS topic to trigger an AWS Lambda function that processes each bid as soon as a user submits it.
F. Refactor the web application to post each incoming bid to an Amazon SQS FIFO queue in place of Kinesis Data Stream
G. Refactor the bid processor to continuously the SQS queu
H. Place the bid processing EC2 instance in an Auto Scaling group with a minimum and a maximum size of 1.
I. Switch the EC2 instance type from t2.large to a larger general compute instance typ
J. Put the bid processor EC2 instances in an Auto Scaling group that scales out the number of EC2 instances running the bid processor, based on the IncomingRecords metric in Kinesis Data Streams.

Answer: C

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

NEW QUESTION 152
A company with multiple accounts is currently using a configuration that does not meet the following security governance policies
• Prevent ingress from port 22 to any Amazon EC2 instance
• Require billing and application tags for resources
• Encrypt all Amazon EBS volumes
A Solutions Architect wants to provide preventive and detective controls including notifications about a specific resource, if there are policy deviations.
Which solution should the Solutions Architect implement?

A. Create an AWS CodeCommit repository containing policy-compliant AWS Cloud Formation templates.Create an AWS Service Catalog portfolio Import the Cloud Formation templates by attaching the CodeCommit repository to the portfolio Restrict users across all accounts to items from the AWSService Catalog portfolio Use AWS Config managed rules to detect deviations from the policie
B. Configure an Amazon CloudWatch Events rule for deviations, and associate a CloudWatch alarm to send notifications when the TriggeredRules metric is greater than zero.
C. Use AWS Service Catalog to build a portfolio with products that are in compliance with the governance policies in a central account Restrict users across all accounts lo AWS Service Catalog products Share a compliant portfolio to other accounts Use AWS Config managed rules to detect deviations from the policies Configure an Amazon CloudWatch Events rule to send a notification when a deviation occurs
D. Implement policy-compliant AWS Cloud Formation templates for each account and ensure that all provisioning is completed by Cloud Formation Configure Amazon Inspector to perform regular checks against resources Perform policy validation and write the assessment output to Amazon CloudWatch Log
E. Create a CloudWatch Logs metric filter to increment a metric when a deviation occurs Configure a CloudWatch alarm to send notifications when the configured metric is greater than zero
F. Restrict users and enforce least privilege access using AWS I A
G. Consolidate all AWS CloudTrail logs into a single account Send the CloudTrail logs to Amazon Elasticsearch Service (Amazon ES). Implement monitoring alerting, and reporting using the Kibana dashboard in Amazon ES and with Amazon SNS.

Answer: C

NEW QUESTION 154
A company that is new to AWS reports it has exhausted its service limits across several accounts that are on the Basic Support plan. The company would like to prevent this from happening in the future.
What is the MOST efficient way of monitoring and managing all service limits in the company's accounts?

A. Use Amazon CloudWatch and AWS Lambda to periodically calculate the limits across all linked accounts using AWS Trusted Advisor, provide notifications using Amazon SNS if the limits are close to exceeding the threshold.
B. Reach out to AWS Support to proactively increase the limits across all account
C. That way, the customer avoids creating and managing infrastructure just to raise the service limits.
D. Use Amazon CloudWatch and AWS Lambda to periodically calculate the limits across all linked accounts using AWS Trusted Advisor, programmatically increase the limits that are close to exceeding the threshold.
E. Use Amazon CloudWatch and AWS Lambda to periodically calculate the limits across all linked accounts using AWS Trusted Advisor, and use Amazon SNS for notifications if a limit is close to exceeding the threshol
F. Ensure that the accounts are using the AWS Business Support plan at a minimum.

Answer: D

Explanation:
https://github.com/awslabs/aws-limit-monitor https://aws.amazon.com/solutions/limit-monitor/

NEW QUESTION 157
A Solutions Architect is working with a company that is extremely sensitive to its IT costs and wishes to implement controls that will result in a predictable AWS spend each month.
Which combination of steps can help the company control and monitor its monthly AWS usage to achieve a cost that is as close as possible to the target amount? (Choose three.)

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

Answer: AEF

NEW QUESTION 158
A company is running multiple applications on Amazon EC2. Each application is deployed and managed by multiple business units. All applications are deployed on a single AWS account but on different virtual private clouds (VPCs). The company uses a separate VPC in the same account for test and development purposes.
Production applications suffered multiple outages when users accidentally terminated and modified resources that belonged to another business unit. A Solutions Architect has been asked to improve the availability of the company applications while allowing the Developers access to the resources they need.
Which option meets the requirements with the LEAST disruption?

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

Answer: C

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

NEW QUESTION 163
A financial company is using a high-performance compute cluster running on Amazon EC2 instances to perform market simulations A DNS record must be created in an Amazon Route 53 private hosted zone when instances start The DNS record must be removed after instances are terminated.
Currently the company uses a combination of Amazon CtoudWatch Events and AWS Lambda to create the
DNS record. The solution worked well in testing with small clusters, but in production with clusters containing thousands of instances the company sees the following error in the Lambda logs:
HTTP 400 error (Bad request).
The response header also includes a status code element with a value of "Throttling" and a status message element with a value of "Rate exceeded "
Which combination of steps should the Solutions Architect take to resolve these issues? (Select THREE)

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

Answer: BEF

NEW QUESTION 164
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your SAP-C01 Exam with Our Prep Materials Via below:

https://www.certleader.com/SAP-C01-dumps.html