Amazon-Web-Services

Exam Questions SAP-C01

AWS Certified Solutions Architect- Professional

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 Solutions Architect must design a highly available, stateless, REST service. The service will require multiple persistent storage layers for service object meta information and the delivery of content. Each request needs to be authenticated and securely processed. There is a requirement to keep costs as low as possible? How can these requirements be met?

A. Use AWS Fargate to host a container that runs a self-contained REST servic
B. Set up an Amazon ECS service that is fronted by an Application Load Balancer (ALB). Use a custom authenticator to control access to the AP
C. Store request meta information in Amazon DynamoDB with Auto Scaling and static content in a secured S3 bucke
D. Make secure signed requests for Amazon S3 objects and proxy the data through the REST service interface.
E. Use AWS Fargate to host a container that runs a self-contained REST servic
F. Set up an ECS service that is fronted by a cross-zone AL
G. Use an Amazon Cognito user pool to control access to the AP
H. Store request meta information in DynamoDB with Auto Scaling and static content in a secured S3 bucke
I. Generate presigned URLs when returning references to content stored in Amazon S3.
J. Set up Amazon API Gateway and create the required API resources and method
K. Use an Amazon Cognito user pool to control access to the AP
L. Configure the methods to use AWS Lambda proxy integrations, and process each resource with a unique AWS Lambda functio
M. Store request meta information in DynamoDB with Auto Scaling and static content in a secured S3 bucke
N. Generate presigned URLs when returning references to content stored in Amazon S3.
O. Set up Amazon API Gateway and create the required API resources and method
P. Use an Amazon API Gateway custom authorizer to control access to the AP
Q. Configure the methods to use AWS Lambda custom integrations, and process each resource with a unique Lambda functio
R. Store request meta information in an Amazon ElastiCache Multi-AZ cluster and static content in a secured S3 bucke
S. Generate presigned URLs when returning references to content stored in Amazon S3.

Answer: C

NEW QUESTION 3
A company receives clickstream data files to Amazon S3 every five minutes. A Python script runs as a cron job once a day on an Amazon EC2 instance to process each file and load it into a database hosted on Amazon RDS. The cron job takes 15 to 30 minutes to process 24 hours of data. The data consumers ask for the data be available as soon as possible.
Which solution would accomplish the desired outcome?

A. Increase the size of the instance to speed up processing and update the schedule to run once an hour.
B. Convert the cron job to an AWS Lambda function and trigger this new function using a cron job on an EC2 instance.
C. Convert the cron job to an AWS Lambda function and schedule it to run once an hour using Amazon CloudWatch events.
D. Create an AWS Lambda function that runs when a file is delivered to Amazon S3 using S3 event notifications.

Answer: D

Explanation:
https://docs.aws.amazon.com/lambda/latest/dg/with-s3.html

NEW QUESTION 4
A large company has many business units. Each business unit has multiple AWS accounts for different purposes. The CIO of the company sees that each business unit has data that would be useful to share with other parts of the company. In total, there are about 10 PB of data that needs to be shared with users in 1,000 AWS accounts. The data is proprietary, so some of it should only be available to users with specific job types. Some of the data is used for throughput of intensive workloads, such as simulations. The number of accounts changes frequently because of new initiatives, acquisitions, and divestitures.
A Solutions Architect has been asked to design a system that will allow for sharing data for use in AWS with all of the employees in the company.
Which approach will allow for secure data sharing in scalable way?

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

Answer: B

NEW QUESTION 5
A company has an Amazon EC2 deployment that has the following architecture:

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

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

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

Answer: B

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

NEW QUESTION 6
A company wants to follow its website on AWS using serverless architecture design patterns for global customers. The company has outlined its requirements as follow:

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

How can the design requirements be met?

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

Answer: C

NEW QUESTION 7
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 8
A company runs a memory-intensive analytics application using on-demand Amazon EC2 compute optimized instance. The application is used continuously and application demand doubles during working hours. The application currently scales based on CPU usage. When scaling in occurs, a lifecycle hook is used because the instance requires 4 minutes to clean the application state before terminating.
Because users reported poor performance during working hours, scheduled scaling actions were implemented so additional instances would be added during working hours. The Solutions Architect has been asked to reduce the cost of the application.
Which solution is MOST cost-effective?

A. Use the existing launch configuration that uses C5 instances, and update the application AMI to include the Amazon CloudWatch agen
B. Change the Auto Scaling policies to scale based on memory utilizatio
C. Use Reserved Instances for the number of instances required after working hours, and use Spot Instances to cover the increased demand during working hours.
D. Update the existing launch configuration to use R5 instances, and update the application AMI to includeSSM Agen
E. Change the Auto Scaling policies to scale based on memory utilizatio
F. Use Reserved instances for the number of instances required after working hours, and use Spot Instances withon-Demand instances to cover the increased demand during working hours.
G. Use the existing launch configuration that uses C5 instances, and update the application AMI to include SSM Agen
H. Leave the Auto Scaling policies to scale based on CPU utilizatio
I. Use scheduled Reserved Instances for the number of instances required after working hours, and use Spot Instances to cover the increased demand during work hours.
J. Create a new launch configuration using R5 instances, and update the application AMI to include the Amazon CloudWatch agen
K. Change the Auto Scaling policies to scale based on memory utilizatio
L. use Reserved Instances for the number of instances required after working hours, and use Standard Reserved Instances with On-Demand Instances to cover the increased demand during working hours.

Answer: D

Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring_ec2.html

NEW QUESTION 9
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 10
A company uses an Amazon EMR cluster to process data once a day. The raw data comes from Amazon S3, and the resulting processed data is also stored in Amazon S3. The processing must complete within 4 hours; currently, it only takes 3 hours. However, the processing time is taking 5 to 10 minutes. longer each week due to an increasing volume of raw data.
The team is also concerned about rising costs as the compute capacity increases. The EMR cluster is currently running on three m3.xlarge instances (one master and two core nodes).
Which of the following solutions will reduce costs related to the increasing compute needs?

A. Add additional task nodes, but have the team purchase an all-upfront convertible Reserved Instance for each additional nod e to offset the costs.

B. Add additional task nodes, but use instance fleets with the master node in on-Demand mode and a mix of On-Demand and Spot Instances for the core and task node
C. Purchase a scheduled Reserved Instances for the master node.
D. Add additional task nodes, but use instance fleets with the master node in Spot mode and a mix of On-Demand and Spot Instances for the core and task node
E. Purchase enough scheduled Reserved Instances to offset the cost of running any On-Demand instances.
F. Add additional task nodes, but use instance fleets with the master node in On-Demand mode and a mix of On-Demand and Spot Instances for the core and task node
G. Purchase a standard all-upfront Reserved Instance for the master node.

Answer: B

NEW QUESTION 10
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 11
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 16
A company runs its containerized batch jobs on Amazon ECS. The jobs are scheduled by submitting a container image, a task definition, and the relevant data to an Amazon S3 bucket. Container images may be unique per job. Running the jobs as quickly as possible is of utmost importance, so submitting jobs artifacts to the S3 bucket triggers the job to run immediately. Sometimes there may no jobs running at all. However, jobs of any size can be submitted with no prior warning to the IT Operations team. Job definitions include CPU and memory resource requirements.
What solution will allow the batch jobs to complete as quickly as possible after being scheduled?

A. Schedule the jobs on an Amazon ECS cluster using the Amazon EC2 launch typ
B. Use Service Auto Scaling to increase or decrease the number of running tasks to suit the number of running jobs.
C. Schedule the jobs directly on EC2 instance
D. Use Reserved Instances for the baseline minimum load, and use On-Demand Instances in an Auto Scaling group to scale up the platform based on demand.
E. Schedule the jobs on an Amazon ECS cluster using the Fargate launch typ
F. Use Service Auto Scaling to increase or decrease the number of running tasks to suit the number of running jobs.
G. Schedule the jobs on an Amazon ECS cluster using the Fargate launch typ
H. Use Spot Instances in an Auto Scaling group to scale the platform based on deman
I. Use Service Auto Scaling to increase or decrease the number of running tasks to suit the number of running jobs.

Answer: C

NEW QUESTION 17
A company has a 24 TB MySQL database in its on-premises data center that grows at the rate of 10 GB per day. The data center is connected to the company's

AWS infrastructure with a 50 Mbps VPN connection.
The company is migrating the application and workload to AWS. The application code is already installed and tested on Amazon EC2. The company now needs to migrate the database and wants to go live on AWS within 3 weeks.
Which of the following approaches meets the schedule with LEAST downtime?

A. 1. Use the VM Import/Export service to import a snapshot on the on-premises database into AWS.2.Launch a new EC2 instance from the snapshot.3. Set up ongoing database replication from on premises to the EC2 database over the VPN.4. Change the DNS entry to point to the EC2 database.5. Stop the replication.
B. 1. Launch an AWS DMS instance.2. Launch an Amazon RDS Aurora MySQL DB instance.3. Configure the AWS DMS instance with on-premises and Amazon RDS database information.4. Start the replication task within AWS DMS over the VPN.5. Change the DNS entry to point to the Amazon RDS MySQL database.6. Stop the replication.
C. 1. Create a database export locally using database-native tools.2. Import that into AWS using AWS Snowball.3. Launch an Amazon RDS Aurora DB instance.4. Load the data in the RDS Aurora DB instance from the export.5. Set up database replication from the on-premises database to the RDS Aurora DB instance over the VPN.6. Change the DNS entry to point to the RDS Aurora DB instance.7. Stop the replication.
D. 1. Take the on-premises application offline.2. Create a database export locally using database-native tools.3. Import that into AWS using AWS Snowball.4. Launch an Amazon RDS Aurora DB instance.5. Load the data in the RDS Aurora DB instance from the export.6. Change the DNS entry to point to the Amazon RDS Aurora DB instance.7. Put the Amazon EC2 hosted application online.

Answer: C

NEW QUESTION 18
A large company has increased its utilization of AWS over time in an unmanaged way. As such, they have a large number of independent AWS accounts across different business units, projects, and environments. The company has created a Cloud Center of Excellence team, which is responsible for managing all aspects of the AWS Cloud, including their AWS accounts.
Which of the following should the Cloud Center of Excellence team do to BEST address their requirements in a centralized way? (Select two.)
A. Control all AWS account root user credential
B. Assign AWS IAM users in the account of each user who needs to access AWS resource
C. Follow the policy of least privilege in assigning permissions to each user.
D. Tag all AWS resources with details about the business unit, project, and environmen
E. Send all AWS Cost and Usage reports to a central Amazon S3 bucket, and use tools such as Amazon Athena and Amazon QuickSight to collect billing details by business unit.
F. Use the AWS Marketplace to choose and deploy a Cost Management too
G. Tag all AWS resources with details about the business unit, project, and environmen
H. Send all AWS Cost and Usage reports for the AWS accounts to this tool for analysis.
I. Set up AWS Organization
J. Enable consolidated billing, and link all existing AWS accounts to a master billing accoun
K. Tag all AWS resources with details about the business unit, project and environmen
L. Analyze Cost and Usage reports using tools such as Amazon Athena and Amazon QuickSight to collect billing details by business unit.
M. Using a master AWS account, create IAM users within the master accoun
N. Define IAM roles in the other AWS accounts, which cover each of the required functions in the accoun
O. Follow the policy of least privilege in assigning permissions to each role, then enable the IAM users to assume the roles that they need to use.

Answer: DE

NEW QUESTION 20
A company is using AWS CloudFormation to deploy its infrastructure. The company is concerned that, if a production CloudFormation stack is deleted, important data stored in Amazon RDS databases or Amazon EBS volumes might also be deleted.
How can the company prevent users from accidentally deleting data in this way?

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

Answer: A

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

NEW QUESTION 24
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 27
A Solutions Architect is responsible for redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application runs on a single high-memory Amazon EC2 instance. It accepts HTTP requests from upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average., most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel.
Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backing. In addition, the current system has issues with availability and data if the single application node fails.
Clients that access this service cannot be modified. They expect to receive a response to each HTTP request they send within 10 seconds before they will time out and retry the request.
Which approach would improve the availability and durability of the system while decreasing the processing latency and minimizing costs?

A. Create an Amazon API Gateway REST API that uses Lambda proxy integration to pass requests to an AWS Lambda functio
B. Migrate the core processing code to a Lambda function and write a wrapper class that provides a handler method that converts the proxy events to the internal application data model and invokes the processing module.
C. Create an Amazon API Gateway REST API that uses a service proxy to put items in an Amazon SQS queu
D. Extract the core processing code from the existing application and update it to pull items from Amazon SQS queu
E. Extract the core processing code from the existing application and update it to pull items from Amazon SQS instead of an in-memory queu
F. Deploy the new processing application to smaller EC2 instances within an Auto Scaling group that scales dynamically based on the approximate number of messages in the Amazon SQS queue.
G. Modify the application to use Amazon DynamoDB instead of Amazon RD
H. Configure Auto Scaling for the DynamoDB tabl
I. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilizatio
J. Back the in-memory queue with a memory-mapped file to an instance store volume and periodically write that file to Amazon S3.
K. Update the application to use a Redis task queue instead of the in-memory queu
L. Build a Docker container image for the applicatio
M. Create an Amazon ECS task definition that includes the application container and a separate container to host Redi
N. Deploy the new task definition as an ECS service using AWS Fargate and enable Auto Scaling.

Answer: B

NEW QUESTION 30
A company is operating a large customer service call center, and stores and processes call recordings with a custom application Approximately 2% of the call recording are transcribed by an offshore team for quality assurance purposes. These recordings take days. The company uses Linux servers for processing the call recording and managing the transcription queue. There is also a web application for the quality assurance staff to review and score call recordings.
The company plans to migrate the system to AWS to reduce storage costs and the time required to transcribe calls.
Which set of actions should be taken to meet the company's objectives?

A. Upload the call recording to Amazon S3 from the call cente
B. Set up an S3 lifecycle policy to move the call recordings to Amazon S3 Glacier after 90 day
C. Use an AWS Lambda trigger to transcribe the call recordings with Amazon Transcrib
D. Use Amazon S3, Amazon API Gateway and Lambda to host the review and scoring application.
E. Upload the call recordings to Amazon S3 from the call cente
F. Set up an S3 lifecycle policy to move the call recordings to Amazon S3 Glacier after 90 day
G. Use an AWS Lambda trigger to transcribe the call recordings with Amazon Mechanical trun
H. Use Amazon EC2 instances in an Auto Scaling group behind an Application Balancer to host the review and scoring application.
I. Use Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer to host the review and scoring applicatio

J. Upload the call recordings to this application from the call center and store them on an Amazon EFS mount poin
K. Use AWS Backup to archive the call recording after 90 day
L. Transcribe the call recordings with Amazon Transcribe.
M. Upload the call recording to Amazon S3 from the call center and put the object key in an Amazon SQS queu
N. Set up an S3 lifecycle policy to move the call recordings to Amazon S3 Glacier after 90 day
O. Use Amazon EC2 instances in the queue as the scaling metri
P. Use Amazon S3, Amazon API Gateway, and AWS Lambda to host the review and scoring application.

Answer: B

NEW QUESTION 33
A company has developed a new billing application that will be released in two weeks. Developers are testing the application running on 10 EC2 instances managed by an Auto Scaling group in subnet 172.31.0.0/24 within VPC A with CIDR block 172.31.0.0/16. The Developers noticed connection timeout errors in the application logs while connecting to an Oracle database running on an Amazon EC2 instance in the same region within VPC B with CIDR block 172.50.0.0/16. The IP of the database instance is hard-coded in the application instances.
Which recommendations should a Solutions Architect present to the Developers to solve the problem in a secure way with minimal maintenance and overhead?

A. Disable the SrcDestCheck attribute for all instances running the application and Oracle Database.Change the default route of VPC A to point ENI of the Oracle Database that has an IP address assigned within the range of 172.50.0.0/26
B. Create and attach internet gateways for both VPC
C. Configure default routes to the Internet gateways for both VPC
D. Assign an Elastic IP for each Amazon EC2 instance in VPC A
E. Create a VPC peering connection between the two VPCs and add a route to the routing table of VPC A that points to the IP address range of 172.50.0.0/16
F. Create an additional Amazon EC2 instance for each VPC as a customer gateway; create one virtual private gateway (VGW) for each VPC, configure an end-to-end VPC, and advertise the routes for 172.50.0.0/16

Answer: C

NEW QUESTION 34
A company is running an email application across multiple AWS Regions. The company uses Ohio (us-east-2) as the primary Region and Northern Virginia (us-east-1) as the Disaster Recovery (DR) Region. The data is continuously replicated from the primary Region to the DR Region by a single instance on the public subnet in both Regions. The replication messages between the Regions have a significant backlog during certain times of the day. The backlog clears on its own after a short time, but it affects the application's RPO.
Which of the following solutions should help remediate this performance problem? (Select TWO)

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

Answer: AC

NEW QUESTION 37
A bank is designing an online customer service portal where customers can chat with customer service agents. The portal is required to maintain a 15-minute RPO or RTO in case of a regional disaster. Banking regulations require that all customer service chat transcripts must be preserved on durable storage for at least 7 years, chat conversations must be encrypted in-flight, and transcripts must be encrypted at rest. The Data Lost Prevention team requires that data at rest must be encrypted using a key that the team controls, rotates, and revokes.
Which design meets these requirements?

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

Answer: B

NEW QUESTION 39
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: B

Explanation:
Import/Export supports importing and exporting data into and out of Amazon S3 buckets. For significant data sets, AWS Import/Export is often faster than Internet transfer and more cost effective than upgrading your connectivity.

NEW QUESTION 44
A company runs a public-facing application that uses a Java-based web sen/ice via a RESTful API It is hosted on Apache Tomcat on a single server in a data center that runs consistently at 30% CPU utilization Use of the API is expected to increase by 10 times with a new product launch The business wants to migrate the application to AWS with no disruption and needs it to scale to meet demand
The company has already decided to use Amazon Route 53 and CNAME records lo redirect traffic How can these requirements be met with the LEAST amount of effort?

A. Use AWS Elastic Beanstalk to deploy the Java web service and enable Auto Scaling Then switch the application to use the new web service
B. Lift and shift the Apache server to the cloud using AWS SMS Then switch the application to direct web service traffic to the new instance
C. Create a Docker image and migrate the image to Amazon ECS Then change the application code to direct web service queries to the ECS container
D. Modify the application to call the web service via Amazon API Gateway Then create a new AWS Lambda Java function to run the Java web service code After testing change API Gateway to use the Lambda function

Answer: A

NEW QUESTION 45
A company is running a .NET three-tier web application on AWS. The team currently uses XL storage optimized instances to store serve the website's image and video files on local instance storage. The company has encountered issues with data loss from replication and instance failures. The Solutions Architect has been asked to redesign this application to improve its reliability while keeping costs low.
Which solution will meet these requirements?

A. Set up a new Amazon EFS share, move all image and video files to this share, and then attach this new drive as a mount point to all existing server
B. Create an Elastic Load Balancer with Auto Scaling general purpose instance
C. Enable Amazon CloudFront to the Elastic Load Balance
D. Enable Cost Explorer and use AWS Trusted advisor checks to continue monitoring the environment for future savings.
E. Implement Auto Scaling with general purpose instance types and an Elastic Load Balance
F. Enable an Amazon CloudFront distribution to Amazon S3 and move images and video files to Amazon S3. Reserve general purpose instances to meet base performance requirement
G. Use Cost Explorer and AWSTrusted Advisor checks to continue monitoring the environment for future savings.
H. Move the entire website to Amazon S3 using the S3 website hosting featur
I. Remove all the web servers and have Amazon S3 communicate directly with the application servers in Amazon VPC.
J. Use AWS Elastic Beanstalk to deploy the .NET applicatio
K. Move all images and video files to Amazon EF
L. Create an Amazon CloudFront distribution that points to the EFS shar
M. Reserve the m4.4xl instances needed to meet base performance requirements.

Answer: B

NEW QUESTION 49
A company stores sales transaction data in Amazon DynamoDB tables. To detect anomalous behaviors and respond quickly, all changes to the items stored in the DynamoDB tables must be logged within 30 minutes. Which solution meets the requirements?

A. Copy the DynamoDB tables into Apache Hive tables on Amazon EMR every hour and analyze them for anomalous behavior
B. Send Amazon SNS notifications when anomalous behaviors are detected.
C. Use AWS CloudTrail to capture all the APIs that change the DynamoDB table
D. Send SNS notifications when anomalous behaviors are detected using CloudTrail event filtering.
E. Use Amazon DynamoDB Streams to capture and send updates to AWS Lambd
F. Create a Lambda function to output records to Amazon Kinesis Data Stream
G. Analyze any anomalies with Amazon Kinesis Data Analytic
H. Send SNS notifications when anomalous behaviors are detected.
I. Use event patterns in Amazon CloudWatch Events to capture DynamoDB API call events with an AWS Lambda function as a target to analyze behavio
J. Send SNS notifications when anomalous behaviors are detected.

Answer: C

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

NEW QUESTION 53
A company is migrating an application to AWS. It wants to use fully managed services as much as possible during the migration. The company needs to store

large, important documents within the application with the following requirements:

- The data must be highly durable and available.
- The data must always be encrypted at rest and in transit.
- The encryption key must be managed by the company and rotated periodically. Which of the following solutions should the Solutions Architect recommend?

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

Answer: B

Explanation:
https://aws.amazon.com/blogs/security/how-to-use-bucket-policies-and-apply-defense-in-depth-to-help-secure-y

NEW QUESTION 54
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 55
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 60
A company used Amazon EC2 instances to deploy a web fleet to host a blog site. The EC2 instances are behind an Application Load Balancer (ALB) and are configured in an Auto Scaling group. The web application stores all blog content on an Amazon EFS volume.
The company recently added a feature for bloggers to add video to their posts, attracting 10 times the previous user traffic. At peak times of day, users report buffering and timeout issues while attempting to reach the site or watch videos.
Which is the MOST cost-efficient and scalable deployment that will resolve the issues for users?

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

Answer: C

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

NEW QUESTION 64
A Solutions Architect must update an application environment within AWS Elastic Beanstalk using a blue/green deployment methodology. The Solutions Architect creates an environment that is identical to the existing application environment and deploys the application to the new environment.
What should be done next to complete the update?

A. Redirect to the new environment using Amazon Route 53
B. Select the Swap Environment URLs option
C. Replace the Auto Scaling launch configuration
D. Update the DNS records to point to the green environment

Answer: B

Explanation:
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAMESwap.html

NEW QUESTION 68
A company runs an IoT platform on AWS. IoT sensors in various locations send data to the company's Node.js API servers on Amazon EC2 instances running behind an Application Load Balancer. The data is stored in an Amazon RDS MySQL DB instance that uses a 4 TB General Purpose SSD volume.
The number of sensors the company has deployed in the field has increased over time, and is expected to grow significantly. The API servers are consistently overloaded and RDS metrics show high write latency.
Which of the following steps together will resolve the issues permanently and enable growth as new sensors are provisioned, while keeping this platform cost-efficient? (Choose two.)

A. Resize the MySQL General Purpose SSD storage to 6 TB to improve the volume's IOPS
B. Re-architect the database tier to use Amazon Aurora instead of an RDS MySQL DB instance and add read replicas
C. Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data
D. Use AWS-X-Ray to analyze and debug application issues and add more API servers to match the load
E. Re-architect the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance

Answer: CE

NEW QUESTION 71
A company runs a legacy system on a single m4.2xlarge Amazon EC2 instance with Amazon EBS2 storage. The EC2 instance runs both the web server and a self-managed Oracle database. A snapshot is made of the EBS volume every 12 hours, and an AMI was created from the fully configured EC2 instance.
A recent event that terminated the EC2 instance led to several hours of downtime. The application was successfully launched from the AMI, but the age of the EBS snapshot and the repair of the database resulted in the loss of 8 hours of data. The system was also down for 4 hours while the Systems Operators manually performed these processes.
What architectural changes will minimize downtime and reduce the chance of lost data?

A. Create an Amazon CloudWatch alarm to automatically recover the instanc
B. Create a script that will check and repair the database upon reboo
C. Subscribe the Operations team to the Amazon SNS message generated by the CloudWatch alarm.
D. Run the application on m4.xlarge EC2 instances behind an Elastic Load Balancer/Application Load Balance
E. Run the EC2 instances in an Auto Scaling group across multiple Availability Zones with a minimum instance count of tw
F. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.
G. Run the application on m4.2xlarge EC2 instances behind an Elastic Load Balancer/Application Load Balance
H. Run the EC2 instances in an Auto Scaling group across multiple Availability Zones with aminimum instance count of on
I. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.
J. Increase the web server instance count to two m4.xlarge instances and use Amazon Route 53 round-robin load balancing to spread the loa
K. Enable Route 53 health checks on the web server
L. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.

Answer: B

Explanation:
Ensures that there are at least two EC instances, each of which is in a different AZ. It also ensures that the database spans multiple AZs. Hence this meets all the criteria.
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html

NEW QUESTION 73
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 75
An enterprise runs 103 line-of-business applications on virtual machines in an on-premises data center. Many of the applications are simple PHP, Java, or Ruby web applications, are no longer actively developed, and serve little traffic.
Which approach should be used to migrate these applications to AWS with the LOWEST infrastructure costs ?

A. Deploy the applications to single-instance AWS Elastic Beanstalk environments without a load balancer.
B. Use AWS SMS to create AMIs for each virtual machine and run them in Amazon EC2.
C. Convert each application to a Docker image and deploy to a small Amazon ECS cluster behind an Application Load Balancer.
D. Use VM Import/Export to create AMIs for each virtual machine and run them in single-instance AWS Elastic Beanstalk environments by configuring a custom image.

Answer: A

Explanation:
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-types.html

NEW QUESTION 76
A company has created an account for individual Development teams, resulting in a total of 200 accounts. All accounts have a single virtual private cloud (VPC) in a single region with multiple microservices running in Docker containers that need to communicate with microservices in other accounts. The Security team requirements state that these microservices must not traverse the public internet, and only certain internal services should be allowed to call other individual services. If there is any denied network traffic for a service, the Security team must be notified of any denied requests, including the source IP.
How can connectivity be established between services while meeting the security requirements?

A. Create a VPC peering connection between the VPC
B. Use security groups on the instances to allow traffic from the security group IDs that are permitted to call the microservic
C. Apply network ACLs to and allow traffic from the local VPC and peered VPCs onl
D. Within the task definition in Amazon ECS for each of the microservices, specify a log configuration by using the awslogs drive
E. Within Amazon CloudWatch Logs, create a metric filter and alarm off of the number of HTTP 403 response
F. Create an alarm when the number of messages exceeds a threshold set by the Security team.
G. Ensure that no CIDR ranges are overlapping, and attach a virtual private gateway (VGW) to each VPC.Provision an IPsec tunnel between each VGW and enable route propagation on the route tabl
H. Configure security groups on each service to allow the CIDR ranges of the VPCs on the other account
I. Enable VPC Flow Logs, and use an Amazon CloudWatch Logs subscription filter for rejected traffi
J. Create an IAM role and allow the Security team to call the AssumeRole action for each account.
K. Deploy a transit VPC by using third-party marketplace VPN appliances running on Amazon EC2, dynamically routed VPN connections between the VPN appliance, and the virtual private gateways (VGWs) attached to each VPC within the regio
L. Adjust network ACLs to allow traffic from the local VPC onl
M. Apply security groups to the microservices to allow traffic from the VPN appliances onl
N. Install the awslogs agent on each VPN appliance, and configure logs to forward to Amazon CloudWatch Logs in the security account for the Security team to access.
O. Create a Network Load Balancer (NLB) for each microservic
P. Attach the NLB to a PrivateLink endpoint service and whitelist the accounts that will be consuming this servic
Q. Create an interface endpoint in the consumer VPC and associate a security group that allows only the security group IDs of the services authorized to call the producer servic
R. On the producer services, create security groups for each microservice and allow only the CIDR range the allowed service
S. Create VPC Flow Logs on each VPC to capture rejected traffic that will be delivered to an Amazon CloudWatch Logs grou
T. Create a CloudWatch Logs subscription that streams the log data to a security account.

Answer: D

Explanation:
AWS PrivateLink provides private connectivity between VPCs, AWS services, and on-premises applications, securely on the Amazon network. AWS PrivateLink makes it easy to connect services across different accounts and VPCs to significantly simplify the network architecture. It seems like the next VPC peering.
https://aws.amazon.com/privatelink/

NEW QUESTION 79
A company is creating an account strategy so that they can begin using AWS. The Security team will provide each team with the permissions they need to follow the principle or least privileged access. Teams would like to keep their resources isolated from other groups, and the Finance team would like each team's resource usage separated for billing purposes.
Which account creation process meets these requirements and allows for changes?

A. Create a new AWS Organizations accoun
B. Create groups in Active Directory and assign them to roles in AWS to grant federated acces
C. Require each team to tag their resources, and separate bills based on tag
D. Control access to resources through IAM granting the minimally required privilege.
E. Create individual accounts for each tea
F. Assign the security as the master account, and enable consolidated billing for all other account
G. Create a cross-account role for security to manage accounts, and send logs to a bucket in the security account.

H. Create a new AWS account, and use AWS Service Catalog to provide teams with the required resources.Implement a third-party billing to provide the Finance team with the resource use for each team based on taggin
I. Isolate resources using IAM to avoid account spraw
J. Security will control and monitor logs and permissions.
K. Create a master account for billing using Organizations, and create each team's account from that master accoun
L. Create a security account for logs and cross-account acces
M. Apply service control policies on each account, and grant the Security team cross-account access to all account
N. Security will create IAM policies for each account to maintain least privilege access.

Answer: B

NEW QUESTION 81
A company runs a Windows Server host in a public subnet that is configured to allow a team of administrators to connect over RDP to troubleshoot issues with hosts in a private subnet. The host must be available at all times outside of a scheduled maintenance window, and needs to receive the latest operating system updates within 3 days of release.
What should be done to manage the host with the LEAST amount of administrative effort?

A. Run the host in a single-instance AWS Elastic Beanstalk environmen
B. Configure the environment with a custom AMI to use a hardened machine image from AWS Marketplac
C. Apply system updates with AWS Systems Manager Patch Manager.
D. Run the host on AWS WorkSpace
E. Use Amazon WorkSpaces Application Manager (WAM) to harden the hos
F. Configure Windows automatic updates to occur every 3 days.
G. Run the host in an Auto Scaling group with a minimum and maximum instance count of 1. Use a hardened machine image from AWS Marketplac
H. Apply system updates with AWS Systems Manager Patch Manager.
I. Run the host in AWS OpsWorks Stack
J. Use a Chief recipe to harden the AMI during instance launch.Use an AWS Lambda scheduled event to run the Upgrade Operating System stack command to apply system updates.

Answer: B

NEW QUESTION 86
A company has a High Performance Computing (HPC) cluster in its on-premises data center which runs thousands of jobs in parallel for one week every month, processing petabytes of images. The images are stored on a network file server, which is replicated to a disaster recovery site. The on-premises data center has reached capacity and has started to spread the jobs out over the course of month in order to better utilize the cluster, causing a delay in the job completion.
The company has asked its Solutions Architect to design a cost-effective solution on AWS to scale beyond the current capacity of 5,000 cores and 10 petabytes of data. The solution must require the least amount of management overhead and maintain the current level of durability.
Which solution will meet the company's requirements?

A. Create a container in the Amazon Elastic Container Registry with the executable file for the jo
B. Use Amazon ECS with Spot Fleet in Auto Scaling group
C. Store the raw data in Amazon EBS SC1 volumes and write the output to Amazon S3.
D. Create an Amazon EMR cluster with a combination of On Demand and Reserved Instance Task Nodes that will use Spark to pull data from Amazon S3. Use Amazon DynamoDB to maintain a list of jobs that need to be processed by the Amazon EMR cluster.
E. Store the raw data in Amazon S3, and use AWS Batch with Managed Compute Environments to create Spot Fleet
F. Submit jobs to AWS Batch Job Queues to pull down objects from Amazon S3 onto Amazon EBS volumes for temporary storage to be processed, and then write the results back to Amazon S3.
G. Submit the list of jobs to be processed to an Amazon SQS to queue the jobs that need to be processed.Create a diversified cluster of Amazon EC2 worker instances using Spot Fleet that will automatically scale based on the queue dept
H. Use Amazon EFS to store all the data sharing it across all instances in the cluster.

Answer: B

NEW QUESTION 91
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 95
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 98
AnyCompany has acquired numerous companies over the past few years. The CIO for AnyCompany would like to keep the resources for each acquired company separate. The CIO also would like to enforce a chargeback model where each company pays for the AWS services it uses.
The Solutions Architect is tasked with designing an AWS architecture that allows AnyCompany to achieve the following:

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

Which of the following approaches would meet these requirements? (Choose two.)

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

Answer: AD

NEW QUESTION 103
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 107
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 111
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 113
The company Security team requires that all data uploaded into an Amazon S3 bucket must be encrypted. The encryption keys must be highly available and the company must be able to control access on a per-user basis, with different users having access to different encryption keys.
Which of the following architectures will meet these requirements? (Choose two.)

A. Use Amazon S3 server-side encryption with Amazon S3-managed key
B. Allow Amazon S3 to generate an AWS/S3 master key, and use IAM to control access to the data keys that are generated.
C. Use Amazon S3 server-side encryption with AWS KMS-managed keys, create multiple customer master keys, and use key policies to control access to them.
D. Use Amazon S3 server-side encryption with customer-managed keys, and use AWS CloudHSM to manage the key
E. Use CloudHSM client software to control access to the keys that are generated.
F. Use Amazon S3 server-side encryption with customer-managed keys, and use two AWS CloudHSM instances configured in high-availability mode to manage the key
G. Use the Cloud HSM client software to control access to the keys that are generated.
H. Use Amazon S3 server-side encryption with customer-managed keys, and use two AWS CloudHSM instances configured in high-availability mode to manage the key
I. Use IAM to control access to the keys that are generated in CloudHSM.

Answer: BD

Explanation:
http://websecuritypatterns.com/blogs/2018/03/01/encryption-and-key-management-in-aws-kms-vs-cloudhsm-mys/

NEW QUESTION 116
An internal security audit of AWS resources within a company found that a number of Amazon EC2 instances running Microsoft Windows workloads were missing several important operating system-level patches. A Solutions Architect has been asked to fix existing patch deficiencies, and to develop a workflow to ensure that future patching requirements are identified and taken care of quickly. The Solutions Architect has decided to use AWS Systems Manager. It is important that EC2 instance reboots do not occur at the same time on all Windows workloads to meet organizational uptime requirements.
Which workflow will meet these requirements in an automated manner?

A. Add a Patch Group tag with a value of Windows Servers to all existing EC2 instance
B. Ensure that all Windows EC2 instances are assigned this ta
C. Associate the AWS-DefaultPatchBaseline to the Windows servers patch grou
D. Define an AWS Systems Manager maintenance window, conduct patching within it, and associate it with the Windows Servers patch grou
E. Register instances with the maintenance window using associated subnet ID
F. Assign the AWS-RunPatchBaseline document as a task within each maintenance window.
G. Add a Patch Group tag a value of Windows Servers to all existing EC2 instance
H. Ensure that all Windows EC2 instances are assigned this ta
I. Associate the AWS-WindowsPatchBaseline document as a task associated with the Windows Servers patch grou
J. Create an Amazon CloudWatch Events rule configured to use a cron expression to schedule the execution of patching using the AWS Systems Manager run comman
K. Create an AWS Systems Manager State Manager document to define commands to be executed during patch execution.
L. Add a Patch Group tag with a value of either Windows Servers1 or Windows Server2 to all existing EC2 instance
M. Ensure that all Windows EC2 instances are assigned this ta
N. Associate theAWS-DefaultPatchBaseline with both Windows Servers patch group
O. Define two non-overlappingAWS Systems Manager maintenance windows, conduct patching within them, and associate each with a different patch grou
P. Register targets with specific maintenance windows using the Patch Group tag
Q. Assign the AWS-RunPatchBaseline document as a task within each maintenance window.
R. Add a Patch Group tag with a value of either Windows servers1 or Windows Server2 to all existing EC2 instance
S. Ensure that all Windows EC2 instances are assigned this ta
T. Associate theAWS-WindowsPatchBaseline with both Windows Servers patch group

. Define two non-overlappingAWS Systems Manager maintenance windows, conduct patching within them, and associate each with a different patch grou
. Assign the AWS-RunWindowsPatchBaseline document as a task within each maintenance windo
. Create an AWS Systems Manager State Manager document to define commands to be executed during patch execution.

Answer: C

NEW QUESTION 120
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 123
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 125
A Solutions Architect is designing the storage layer for a data warehousing application. The data files are large, but they have statically placed metadata at the beginning of each file that describes the size and placement of the file's index. The data files are read in by a fleet of Amazon EC2 instances that store the index size, index location, and other category information about the data file in a database. That database is used by Amazon EMR to group files together for deeper analysis.
What would be the MOST cost-effective, high availability storage solution for this workflow?

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

Answer: A

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

NEW QUESTION 128
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 133
A company has a website that enables users to upload videos Company policy states the uploaded videos must be analyzed for restricted content An uploaded video is placed in Amazon S3, and a message is pushed to an Amazon SQS queue with the video's location A backend application pulls this location from Amazon SQS and analyzes the video
The video analysis is compute-intensive and occurs sporadically during the day The website scales with demand The video analysis application runs on a fixed number of instances Peak demand occurs during the holidays, so the company must add instances to the application during this time All instances used are currently on-demand Amazon EC2 T2 instances The company wants to reduce the cost of the current solution.
Which of the following solutions is MOST cost-effective?

A. Keep the website on T2 instances Determine the minimum number of website instances required during off-peak times and use Spot Instances to cover them while using Reserved Instances to covet peak demand Use Amazon EC2 R4 and Amazon EC2 R5 Reserved Instances in an Auto Scaling group for the video

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

Answer: B

NEW QUESTION 134
An organization has a write-intensive mobile application that uses Amazon API Gateway, AWS Lambda, and Amazon DynamoDB. The application has scaled well, however, costs have increased exponentially because of higher than anticipated Lambda costs. The application's use is unpredictable, but there has been a steady 20% increase in utilization every month.
While monitoring the current Lambda functions, the Solutions Architect notices that the execution-time averages 4.5 minutes. Most of the wait time is the result of a high-latency network call to a 3-TB MySQL database server that is on-premises. A VPN is used to connect to the VPC, so the Lambda functions have been configured with a five-minute timeout.
How can the Solutions Architect reduce the cost of the current architecture?

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

Answer: D

NEW QUESTION 135
A Solution Architect is designing a deployment strategy for an application tier and gas the following requirements.
* The application code will need a 500 HB static dataset to be present before application startup.
* The application tier be able to scale Up and down based on demand with as little startup time as possible.
* The development team should be able to update the code multiple times each day.
* Critical operating system (OS) patches must be installed within 48 hours of being released. Which deployment strategy meets these requirements?

A. Use AWS Manager to create a new AMI with the updated OS patches . Update the Auto Scaling group to use the patches AMI and replace existing unpatche
B. Use AWS CodeDeploy to push the application code to the instance
C. Store the static data in Amazon EFS.
D. Use AWS System Manager to create a new AMI with upload OS patche
E. Update the Auto Scaling group to use the patches AMI and replace existing unpatches and the application code as a batch job every nigh
F. Store the static data in Amazon EFS.
G. Use an Amazon provided AMI for the OS Configure an Auto Scaling group set to a static instance coun
H. Configure an Amazon EC2 data script to download the data from Amazon S3 install OS patches with AWS system Manager when they are release
I. Use Codedeploy to push the application code to the instances.
J. Use an Amazon provided AMI for the OS Configure an Auto Scaling group Configure an Amazon EC2 user data script to download the data from Amazon S3. Replace existing instances after eachAmazon-provided AMI releas
K. Use AWS CodeDeploy to push the application code to the instances.

Answer: C

NEW QUESTION 136
A company is currently using AWS CodeCommit for its source control and AWS CodePipeline for continuous integration. The pipeline has a build stage for building the artifacts which is then staged in an Amazon S3 bucket.
The company has identified various improvement opportunities in the existing process, and a Solutions Architect has been given the following requirement:

- Create a new pipeline to support feature development
- Support feature development without impacting production applications
- Incorporate continuous testing with unit tests
- Isolate development and production artifacts
- Support the capability to merge tested code into production code. How should the Solutions Architect achieve these requirements?

A. Trigger a separate pipeline from CodeCommit feature branche
B. Use AWS CodeBuild for running unit test
C. Use CodeBuild to stage the artifacts within an S3 bucket in a separate testing account.
D. Trigger a separate pipeline from CodeCommit feature branche
E. Use AWS Lambda for running unit test
F. Use AWS CodeDeploy to stage the artifacts within an S3 bucket in a separate testing account.
G. Trigger a separate pipeline from CodeCommit tags Use Jenkins for running unit test
H. Create a stage in the pipeline with S3 as the target for staging the artifacts with an S3 bucket in a separate testing account.

I. Create a separate CodeCommit repository for feature development and use it to trigger the pipelin
J. Use AWS Lambda for running unit test
K. Use AWS CodeBuild to stage the artifacts within different S3 buckets in the same production account.

Answer: A

Explanation:
https://docs.aws.amazon.com/codebuild/latest/userguide/how-to-create-pipeline.html

NEW QUESTION 141
A company has decided to move some workloads onto AWS to create a grid environment to run market analytics. The grid will consist of many similar instances, spun-up by a job-scheduling function. Each time a large analytics workload is completed, a new VPC is deployed along with job scheduler and grid nodes. Multiple grids could be running in parallel.
Key requirements are:

- Grid instances must communicate with Amazon S3 retrieve data to be processed.
- Grid instances must communicate with Amazon DynamoDB to track intermediate data,
- The job scheduler need only to communicate with the Amazon EC2 API to start new grid nodes.

A key requirement is that the environment has no access to the internet, either directly or via the on-premises proxy. However, the application needs to be able to seamlessly communicate to Amazon S3, Amazon DynamoDB, and Amazon EC2 API, without the need for reconfiguration for each new deployment.
Which of the following should the Solutions Architect do to achieve this target architecture? (Choose three.)

A. Enable VPC endpoints for Amazon S3 and DynamoDB.
B. Disable Private DNS Name Support.
C. Configure the application on the grid instances to use the private DNS name of the Amazon S3 endpoint.
D. Populate the on-premises DNS server with the private IP addresses of the EC2 endpoint.
E. Enable an interface VPC endpoint for EC2.
F. Configure Amazon S3 endpoint policy to permit access only from the grid nodes.

Answer: ACE

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/connect-s3-vpc-endpoint/ https://docs.aws.amazon.com/vpc/latest/userguide/vpce-interface.html

NEW QUESTION 145
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 148
A company is finalizing the architecture for its backup solution for applications running on AWS. All of the applications run on AWS and use at least two Availability Zones in each tier.
Company policy requires IT to durably store nightly backups for all its data in at least two locations: production and disaster recovery. The locations must be in different geographic regions. The company also needs the backup to be available to restore immediately at the production data center, and within 24 hours at the disaster recovery location. All backup processes must be fully automated.
What is the MOST cost-effective backup solution that will meet all requirements?

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

I. Back up all the data to Amazon S3 in the production regio
J. Set up cross-region replication of this S3 bucket to another region and set up a lifecycle policy in the second region to immediately move this data to Amazon Glacier.

Answer: D

NEW QUESTION 149
A utility company wants to collect usage data every 5 minutes from its smart meters to facilitate time-of-use metering When a meter sends data to AWS the data is sent to Amazon API Gateway, processed by an AWS Lambda function and stored in an Amazon DynamoDB table During the pilot phase, the Lambda functions took from 3 to 5 seconds to complete
As more smart meters are deployed, the Engineers notice the Lambda functions are taking from 1 to 2 minutes to complete The functions are also increasing in duration as new types of metrics are collected from the devices There are many ProvisionedThroughputExceededException errors while performing PUT operations on DynamoDB and there are also many TooMany Requests Exception errors from Lambda.
Which combination of changes will resolve these issues? (Select TWO)

A. increase the write capacity units to the DynamoDB table
B. Increase the memory available to the Lambda functions
C. Increase the payload size from the smart meters to send more data
D. Stream the data into an Amazon Kinesis data stream from API Gateway and process the data in batches
E. Collect data in an Amazon SQS FIFO queue, which triggers a Lambda function to process each message

Answer: AB

NEW QUESTION 153
A company is running a web application with On-Demand Amazon EC2 instances in Auto Scaling groups that scale dynamically based on custom metrics After extensive testing the company determines that the m5 2xlarge instance size is optimal for the workload Application data is stored in db r4 4xlarge Amazon RDS instances that are confirmed to be optimal The traffic to the web application spikes randomly during the day
What other cost-optimization methods should the company implement to further reduce costs without impacting the reliability of the application?

A. Double the instance count in the Auto Scaling groups and reduce the instance size to m5 large
B. Reserve capacity for the RDS database and the minimum number of EC2 instances that are constantly running
C. Reduce the RDS instance size to db r4 xlarge and add five equivalents sized read replicas to provide reliability
D. Reserve capacity for all EC2 instances and leverage Spot Instance pricing for the RDS database

Answer: B

NEW QUESTION 157
A company currently runs a secure application on Amazon EC2 that takes files from on-premises locations through AWS Direct Connect, processes them, and uploads them to a single Amazon S3 bucket. The application uses HTTPS for encryption in transit to Amazon S3, and S3 server-side encryption to encrypt at rest.
Which of the following changes should the Solutions Architect recommend to make this solution more secure without impeding application's performance?

A. Add a NAT gatewa
B. Update the security groups on the EC2 instance to allow access to and from the S3 IP range onl
C. Configure an S3 bucket policy that allows communication from the NAT gateway's Elastic IP address only.
D. Add a VPC endpoin
E. Configure endpoint policies on the VPC endpoint to allow access to the required Amazon S3 buckets onl
F. Implement an S3 bucket policy that allows communication from the VPC's source IP range only.
G. Add a NAT gatewa
H. Update the security groups on the EC2 instance to allow access to and from the S3 IP range onl
I. Configure an S3 bucket policy that allows communication from the source public IP address of the on-premises network only.
J. Add a VPC endpoin
K. Configure endpoint policies on the VPC endpoint to allow access to the required S3 buckets onl
L. Implement an S3 bucket policy that allows communication from the VPC endpoint only.

Answer: D

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/dev/example-bucket-policies-vpc-endpoint.html

NEW QUESTION 161
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 166
A company is using an Amazon CloudFront distribution to distribute both static and dynamic content from a web application running behind an Application Load Balancer. The web application requires user authorization and session tracking for dynamic content. The CloudFront distribution has a single cache behavior configured to forward the Authorization, Host, and User-Agent HTTP whitelist headers and a session cookie to the origin. All other cache behavior settings are set

to their default value.
A valid ACM certificate is applied to the CloudFront distribution with a matching CNAME in the distribution settings. The ACM certificate is also applied to the HTTPS listener for the Application Load Balancer. The CloudFront origin protocol policy is set to HTTPS only. Analysis of the cache statistics report shows that the miss rate for this distribution is very high.
What can the Solutions Architect do to improve the cache hit rate for this distribution without causing the SSL/TLS handshake between CloudFront and the Application Load Balancer to fail?

A. Create two cache behaviors for static and dynamic conten
B. Remove the User-Agent and Host HTTP headers from the whitelist headers section on both if the cache behavior
C. Remove the session cookie from the whitelist cookies section and the Authorization HTTP header from the whitelist headers section for cache behavior configured for static content.
D. Remove the User-Agent and Authorization HTTP headers from the whitelist headers section of the cache behavio
E. Then update the cache behavior to use presigned cookies for authorization.
F. Remove the Host HTTP header from the whitelist headers section and remove the session cookie from the whitelist cookies section for the default cache behavio
G. Enable automatic object compression and use Lambda@Edge viewer request events for user authorization.
H. Create two cache behaviors for static and dynamic conten
I. Remove the User-Agent HTTP header from the whitelist headers section on both of the cache behavior
J. Remove the session cookie from the whitelist cookies section and the Authorization HTTP header from the whitelist headers section for cache behavior configured for static content.

Answer: D

NEW QUESTION 171
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 175
A company is planning the migration of several lab environments used for software testing. An assortment of custom tooling is used to manage the test runs for each lab. The labs use immutable infrastructure for the software test runs, and the results are stored in a highly available SQL database cluster. Although completely rewriting the custom tooling is out of scope for the migration project, the company would like to optimize workloads during the migration.
Which application migration strategy meets this requirement?

A. Re-host
B. Re-platform
C. Re-factor/re-architect
D. Retire

Answer: B

Explanation:
https://aws.amazon.com/blogs/enterprise-strategy/6-strategies-for-migrating-applications-to-the-cloud/

NEW QUESTION 179
The Security team needs to provide a team of interns with an AWS environment so they can build the serverless video transcoding application. The project will use Amazon S3, AWS Lambda, Amazon API Gateway, Amazon Cognito, Amazon DynamoDB, and Amazon Elastic Transcoder.
The interns should be able to create and configure the necessary resources, but they may not have access to create or modify AWS IAM roles. The Solutions Architect creates a policy and attaches it to the interns' group.
How should the Security team configure the environment to ensure that the interns are self-sufficient?

A. Create a policy that allows creation of project-related resources onl
B. Create roles with required service permissions, which are assumable by the services.
C. Create a policy that allows creation of all project-related resources, including roles that allow access only to specified resources.
D. Create roles with the required service permissions, which are assumable by the service
E. Have theinterns create and use a bastion host to create the project resources in the project subnet only.
F. Create a policy that allows creation of project-related resources onl
G. Require the interns to raise a request for roles to be created with the Security tea
H. The interns will provide the requirements for the permissions to be set in the role.

Answer: A

NEW QUESTION 182
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 186
A company wants to move a web application to AWS. The application stores session information locally on each web server, which will make auto scaling difficult. As part of the migration, the application will be rewritten to decouple the session data from the web servers. The company requires low latency, scalability, and availability.
Which service will meet the requirements for storing the session information in the MOST cost-effective way?

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

Answer: D

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

NEW QUESTION 189
During a security audit of a Service team's application a Solutions Architect discovers that a username and password tor an Amazon RDS database and a set of AWSIAM user credentials can be viewed in the AWS Lambda function code. The Lambda function uses the username and password to run queries on the database and it uses the I AM credentials to call AWS services in a separate management account.
The Solutions Architect is concerned that the credentials could grant inappropriate access to anyone who can view the Lambda code The management account and the Service team's account are in separate AWS Organizations organizational units (OUs)
Which combination of changes should the Solutions Architect make to improve the solution's security? (Select TWO)

A. Configure Lambda to assume a tole in the management account with appropriate access to AWS
B. Configure Lambda to use the stored database credentials in AWS Secrets Manager and enable automatic rotation
C. Create a Lambda function to rotate the credentials every hour by deploying a new Lambda version with the updated credentials
D. Use an SCP on the management accounts OU to prevent IAM users from accessing resources m the Service team's account
E. Enable AWS Shield Advanced on the management account to shield sensitive resources from unauthorized IAM access

Answer: BD

NEW QUESTION 191
A large global company wants to migrate a stateless mission-critical application to AWS. The application is based on IBM WebSphere (application and integration middleware), IBM MQ (messaging middleware), and IBM DB2 (database software) on a z/OS operating system.
How should the Solutions Architect migrate the application to AWS?

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

Answer: B

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

NEW QUESTION 195
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 199
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 201
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 203
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

SAP-C01 Practice Exam Features:

* SAP-C01 Questions and Answers Updated Frequently

* SAP-C01 Practice Questions Verified by Expert Senior Certified Staff

* SAP-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* SAP-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year