DOP-C02 Q&As

AWS Certified DevOps Engineer - Professional

Pass Amazon DOP-C02 Exam with 100% Guarantee

Following Questions and Answers are all new published by Amazon Official Exam Center

QUESTION 1

A company manages multiple AWS accounts in AWS Organizations. The company\\'s security policy states that AWS account root user credentials for member accounts must not be used. The company monitors access to the root user credentials.

A recent alert shows that the root user in a member account launched an Amazon EC2 instance. A DevOps engineer must create an SCP at the organization\\'s root level that will prevent the root user in member accounts from making any AWS service API calls.

Which SCP will meet these requirements?

A.

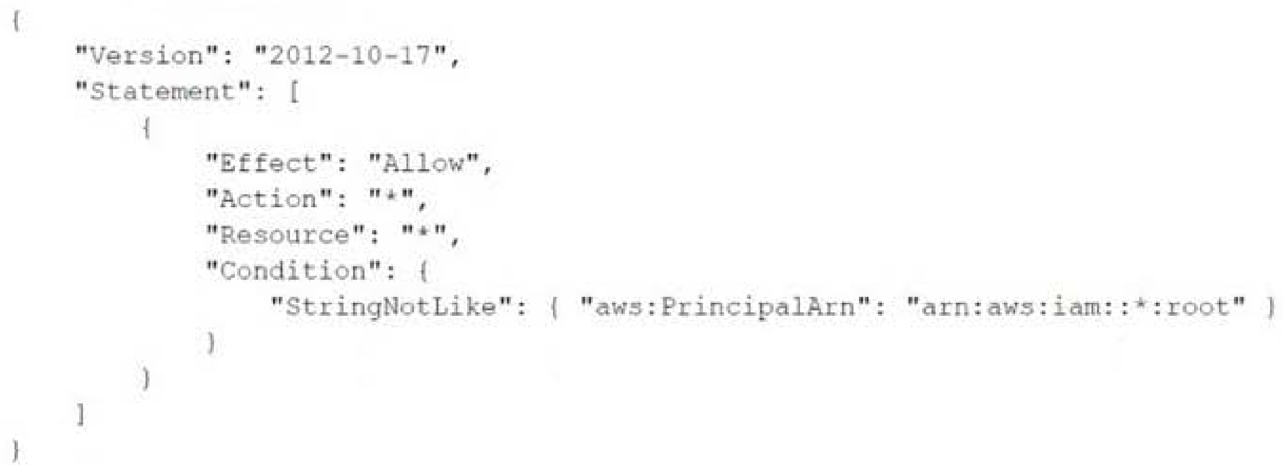

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "*",
            "Resource": "*",
            "Condition": {
                "StringNotLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
            }
        }
    ]
}
```

B.

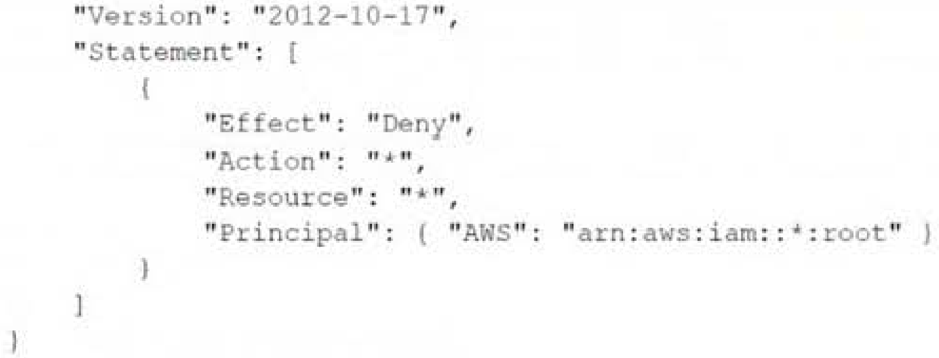

```
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Deny",
            "Action": "*",
            "Resource": "*",
            "Principal": { "AWS": "arn:aws:iam::*:root" }
        }
    ]
}
```

C.

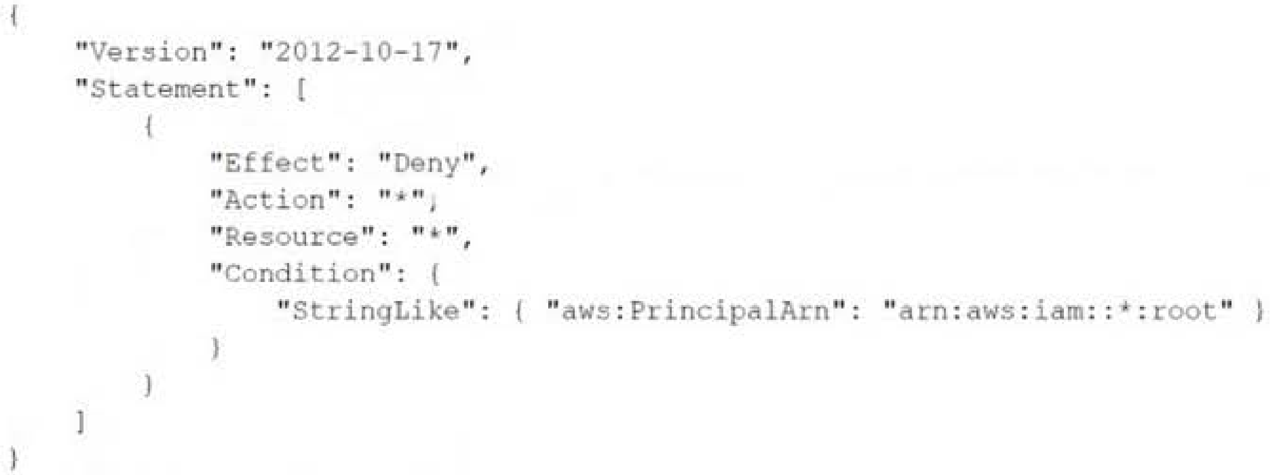

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Deny",
            "Action": "*",
            "Resource": "*",
            "Condition": {
                "StringLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
            }
        }
    ]
}
```

D.

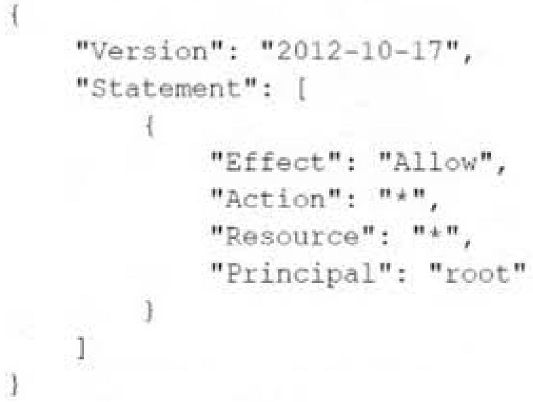

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "*",
            "Resource": "*",
            "Principal": "root"
        }
    ]
}
```

A. Option A

B. Option B

C. Option C

D. Option D

Correct Answer: D

QUESTION 2

An application running on a set of Amazon EC2 instances in an Auto Scaling group requires a configuration file to operate. The instances are created and maintained with AWS CloudFormation. A DevOps engineer wants the instances to have the latest configuration file when launched and wants changes to the configuration file to be reflected on all the instances with a minimal delay when the CloudFormation template is updated. Company policy requires that application configuration files be maintained along with AWS infrastructure configuration files m source control.

Which solution will accomplish this?

A. In the CloudFormaiion template add an AWS Config rule. Place the configuration file content in the rule\\'s InputParameters property and set the Scope property to the EC2 Auto Scaling group. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.

B. In the CloudFormation template add an EC2 launch template resource. Place the configuration file content in the launch template. Configure the cfn-mit script to run when the instance is launched and configure the cfn-hup script to poll for updates to the configuration.

C. In the CloudFormation template add an EC2 launch template resource. Place the configuration file content in the launch template. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.

D. In the CloudFormation template add CloudFormation imt metadata. Place the configuration file content m the metadata. Configure the cfn-init script to run when the instance is launched and configure the cfn-hup script to poll for updates to the configuration.

Correct Answer: D

Use the AWS::CloudFormation::Init type to include metadata on an Amazon EC2 instance for the cfn-init helper script. If your template calls the cfn-init script, the script looks for resource metadata rooted in the AWS::CloudFormation::Init metadata key.

Reference: https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-init.html

QUESTION 3

A company is deploying a new application that uses Amazon EC2 instances. The company needs a solution to query application logs and AWS account API activity Which solution will meet these requirements?

A. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon CloudWatch Logs Configure AWS CloudTrail to deliver the API logs to Amazon S3 Use CloudWatch to query both sets of logs.

B. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon CloudWatch Logs Configure AWS CloudTrail to deliver the API logs to CloudWatch Logs Use CloudWatch Logs Insights to query both sets of logs.

C. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon Kinesis Configure AWS CloudTrail to deliver the API logs to Kinesis Use Kinesis to load the data into Amazon Redshift Use Amazon Redshift to query both sets of logs.

D. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon S3 Use AWS CloudTrail to deliver the API togs to Amazon S3 Use Amazon Athena to query both sets of logs in Amazon S3.

Correct Answer: D

This solution will meet the requirements because it will use Amazon S3 as a common data lake for both the application logs and the API logs. Amazon S3 is a service that provides scalable, durable, and secure object storage for any type of data. You can use the Amazon CloudWatch agent to send logs from your EC2 instances to S3 buckets, and use AWS CloudTrail to deliver the API logs to S3 buckets as well. You can also use Amazon Athena to query both sets of logs in S3 using standard SQL, without loading or transforming them. Athena is a serverless interactive query service that allows you to analyze data in S3 using a variety of data formats, such as JSON, CSV, Parquet, and ORC.

QUESTION 4

On which local address does the Docker DNS server listen?

A. 127.0.0.1

B. 127.0.0.111

C. 127.0.0.254

D. 127.0.0.11

Correct Answer: D

Note: If you need access to a host\\'s localhost resolver, you must modify your DNS service on the host to listen on a non-localhost address that is reachable from within the container. Note: The DNS server is always at 127.0.0.11.

Reference: https://docs.docker.com/engine/userguide/networking/configure-dns/

QUESTION 5

A company is developing an application that will generate log events. The log events consist of five distinct metrics every one tenth of a second and produce a large amount of data The company needs to configure the application to write the logs to Amazon Time stream The company will configure a daily query against the Timestream table.

Which combination of steps will meet these requirements with the FASTEST query performance? (Select THREE.)

A. Use batch writes to write multiple log events in a Single write operation

B. Write each log event as a single write operation

C. Treat each log as a single-measure record

D. Treat each log as a multi-measure record

E. Configure the memory store retention period to be longer than the magnetic store retention period

F. Configure the memory store retention period to be shorter than the magnetic store retention period

Correct Answer: ADF

A comprehensive and detailed explanation is: Option A is correct because using batch writes to write multiple log events in a single write operation is a recommended practice for optimizing the performance and cost of data ingestion in Timestream. Batch writes can reduce the number of network round trips and API calls, and can also take advantage of parallel processing by Timestream. Batch writes can also improve the compression ratio of data in the memory store and the magnetic store, which can reduce the storage costs and improve the query performance1. Option B is incorrect because writing each log event as a single write operation is not a recommended practice for optimizing the performance and cost of data ingestion in Timestream. Writing each log event as a single write operation would increase the number of network round trips and API calls, and would also reduce the compression ratio of data in the memory store and the magnetic store. This would increase the storage costs and degrade the query performance1. Option C is incorrect because treating each log as a single-measure record is not a recommended practice for optimizing the query performance in Timestream. Treating each log as a single-measure record would result in creating multiple records for each timestamp, which would increase the storage size and the query latency. Moreover, treating each log as a single-measure record would require using joins to query multiple measures for the same timestamp, which would add complexity and overhead to the query processing2. Option D is correct because treating each log as a multi-measure record is a recommended practice for optimizing the query performance in Timestream. Treating each log as a multi-measure record would result in creating a single record for each timestamp, which would reduce the storage size and the query latency. Moreover, treating each log as a multi-measure record would allow querying multiple measures for the same timestamp without using joins, which would simplify and speed up the query processing2. Option E is incorrect because configuring the memory store retention period to be longer than the magnetic store retention period is not a valid option in Timestream. The memory store retention period must always be shorter than or equal to the magnetic store retention period. This ensures that data is moved from the memory store to the magnetic store before it expires out of the memory store3. Option F is correct because configuring the memory store retention period to be shorter than the magnetic store retention period is a valid option in Timestream. The memory store retention period determines how long data is kept in the memory store, which is optimized for fast point-in-time queries. The magnetic store retention period determines how long data is kept in the magnetic store, which is optimized for fast analytical queries. By configuring these retention periods appropriately, you can balance your storage costs and query performance according to your application needs3. References:

1: Batch writes

2: Multi-measure records vs. single-measure records

3: Storage

QUESTION 6

A company wants to use a grid system for a proprietary enterprise m-memory data store on top of AWS. This system can run in multiple server nodes in any Linux-based distribution. The system must be able to reconfigure the entire cluster every time a node is added or removed. When adding or removing nodes an /etc./cluster/nodes config file must be updated listing the IP addresses of the current node members of that cluster.

The company wants to automate the task of adding new nodes to a cluster.

What can a DevOps engineer do to meet these requirements?

A. Use AWS OpsWorks Stacks to layer the server nodes of that cluster. Create a Chef recipe that populates the content of the \\'etc./cluster/nodes config file and restarts the service by using the current members of the layer. Assign that recipe to the Configure lifecycle event.

B. Put the file nodes config in version control. Create an AWS CodeDeploy deployment configuration and deployment group based on an Amazon EC2 tag value for thecluster nodes. When adding a new node to the cluster update the file with all tagged instances and make a commit in version control. Deploy the new file and restart the services.

C. Create an Amazon S3 bucket and upload a version of the /etc./cluster/nodes config file Create a crontab script that will poll for that S3 file and download it frequently. Use a process manager such as Monit or system, to restart the cluster services when it detects that the new file was modified. When adding a node to the cluster edit the file\\'s most recent members Upload the new file to the S3 bucket.

D. Create a user data script that lists all members of the current security group of the cluster and automatically updates the /etc/cluster/. nodes config. Tile whenever a new instance is added to the cluster.

Correct Answer: A

You can run custom recipes manually, but the best approach is usually to have AWS OpsWorks Stacks run them automatically. Every layer has a set of built-in recipes assigned each of five lifecycle events--Setup, Configure, Deploy, Undeploy, and Shutdown. Each time an event occurs for an instance, AWS OpsWorks Stacks runs the associated recipes for each of the instance\\'s layers, which handle the corresponding tasks. For example, when an instance finishes booting, AWS OpsWorks Stacks triggers a Setup event. This event runs the associated layer\\'s Setup recipes, which typically handle tasks such as installing and configuring packages

QUESTION 7

A DevOps engineer at a company is supporting an AWS environment in which all users use AWS IAM Identity Center (AWS Single Sign-On). The company wants to immediately disable credentials of any new IAM user and wants the security team to receive a notification.

Which combination of steps should the DevOps engineer take to meet these requirements? (Choose three.)

A. Create an Amazon EventBridge rule that reacts to an IAM CreateUser API call in AWS CloudTrail.

B. Create an Amazon EventBridge rule that reacts to an IAM GetLoginProfile API call in AWS CloudTrail.

C. Create an AWS Lambda function that is a target of the EventBridge rule. Configure the Lambda function to disable any access keys and delete the login profiles that are associated with the IAM user.

D. Create an AWS Lambda function that is a target of the EventBridge rule. Configure the Lambda function to delete the login profiles that are associated with the IAM user.

E. Create an Amazon Simple Notification Service (Amazon SNS) topic that is a target of the EventBridge rule. Subscribe the security team\\'s group email address to the topic.

F. Create an Amazon Simple Queue Service (Amazon SQS) queue that is a target of the Lambda function. Subscribe the security team\\'s group email address to the queue.

Correct Answer: ACE

QUESTION 8

You have an application running a specific process that is critical to the application\\'s functionality, and have added the health check process to your Auto Scaling Group. The instances are showing healthy but the application itself is not working as it should. What could be the issue with the health check, since it is still showing the instances as healthy.

A. You do not have the time range in the health check properly configured

B. It is not possible for a health check to monitor a process that involves the application

C. The health check is not configured properly

D. The health check is not checking the application process

Correct Answer: D

If you have custom health checks, you can send the information from your health checks to Auto Scaling so that Auto Scaling can use this information. For example, if you determine that an instance is not functioning as expected, you can set the health status of the instance to Unhealthy. The next time that Auto Scaling performs a health check on the instance, it will determine that the instance is unhealthy and then launch a replacement instance.

QUESTION 9

The security team depends on AWS CloudTrail to detect sensitive security issues in the company\\'s AWS account. The DevOps engineer needs a solution to auto-remediate CloudTrail being turned off in an AWS account.

What solution ensures the LEAST amount of downtime for the CloudTrail log deliveries?

A. Create an Amazon EventBridge rule for the CloudTrail StopLogging event. Create an AWS Lambda (unction that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was called. Add the Lambda function ARN as a target to the EventBridge rule.

B. Deploy the AWS-managed CloudTrail-enabled AWS Config rule set with a periodic interval to 1 hour. Create an Amazon EventBridge rule tor AWS Config rules compliance change. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLoggmg was called. Add the Lambda function ARN as a target to the EventBridge rule.

C. Create an Amazon EventBridge rule for a scheduled event every 5 minutes. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on a CloudTrail trail in the AWS account. Add the Lambda function ARN as a target to the EventBridge rule.

D. Launch a t2 nano instance with a script running every 5 minutes that uses the AWS SDK to query CloudTrail in the current account. If the CloudTrail trail is disabled have the script re-enable the trail.

Correct Answer: A

https://aws.amazon.com/blogs/mt/monitor-changes-and-auto-enable-logging-in-aws-cloudtrail/

QUESTION 10

A company is using an organization in AWS Organizations to manage multiple AWS accounts. The company\\'s development team wants to use AWS Lambda functions to meet resiliency requirements and is rewriting all applications to work with Lambda functions that are deployed in a VPC. The development team is using Amazon Elastic Pile System (Amazon EFS) as shared storage in Account A in the organization.

The company wants to continue to use Amazon EPS with Lambda Company policy requires all serverless projects to be deployed in Account B.

A DevOps engineer needs to reconfigure an existing EFS file system to allow Lambda functions to access the data through an existing EPS access point.

Which combination of steps should the DevOps engineer take to meet these requirements? (Select THREE.)

A. Update the EFS file system policy to provide Account B with access to mount and write to the EFS file system in Account A.

B. Create SCPs to set permission guardrails with fine-grained control for Amazon EFS.

C. Create a new EFS file system in Account B Use AWS Database Migration Service (AWS DMS) to keep data from Account A and Account B synchronized.

D. Update the Lambda execution roles with permission to access the VPC and the EFS file system.

E. Create a VPC peering connection to connect Account A to Account B.

F. Configure the Lambda functions in Account B to assume an existing IAM role in Account A

Correct Answer: AEF

A Lambda function in one account can mount a file system in a different account. For this scenario, you configure VPC peering between the function VPC and the file system VPC. https://docs.aws.amazon.com/lambda/latest/dg/servicesefs.html https://aws.amazon.com/ru/blogs/storage/mount-amazon-efs-file-systems-cross-account-from-amazon-eks/

1.

Need to update the file system policy on EFS to allow mounting the file system into Account B. ## File System Policy $ cat file-system-policy.json { "Statement": [{ "Effect": "Allow", "Action": ["elasticfilesystem:ClientMount", "elasticfilesystem:ClientWrite"], "Principal": { "AWS": "arn:aws:iam:::root" # Replace with AWS account ID of EKS cluster } }] }

2.

Need VPC peering between Account A and Account B as the pre-requisite

3.

Need to assume cross-account IAM role to describe the mounts so that a specific mount can be chosen.

QUESTION 11

A company is launching an application that stores raw data in an Amazon S3 bucket. Three applications need to access the data to generate reports. The data must be redacted differently for each application before the applications can access the data.

Which solution will meet these requirements?

A. Create an S3 bucket for each application. Configure S3 Same-Region Replication (SRR) from the raw data\\'s S3 bucket to each application\\'s S3 bucket. Configure each application to consume data from its own S3 bucket.

B. Create an Amazon Kinesis data stream. Create an AWS Lambda function that is invoked by object creation events in the raw data\\'s S3 bucket. Program the Lambda function to redact data for each application. Publish the data on the Kinesis data stream. Configure each application to consume data from the Kinesis data stream.

C. For each application, create an S3 access point that uses the raw data\\'s S3 bucket as the destination. Create an AWS Lambda function that is invoked by object creation events in the raw data\\'s S3 bucket. Program the Lambda function to redact data for each application. Store the data in each application\\'s S3 access point. Configure each application to consume data from its own S3 access point.

D. Create an S3 access point that uses the raw data\\'s S3 bucket as the destination. For each application, create an S3 Object Lambda access point that uses the S3 access point. Configure the AWS Lambda function for each S3 Object Lambda access point to redact data when objects are retrieved. Configure each application to consume data from its own S3 Object Lambda access point.

Correct Answer: D

The best solution is to use S3 Object Lambda1, which allows you to add your own code to S3 GET, LIST, and HEAD requests to modify and process data as it is returned to an application2. This way, you can redact the data differently for

each application without creating and storing multiple copies of the data or running proxies.

The other solutions are less efficient or scalable because they require replicating the data to multiple buckets, streaming the data through Kinesis, or storing the data in S3 access points.

References: 1: Amazon S3 Features | Object Lambda | AWS 2: Transforming objects with S3 Object Lambda -Amazon Simple Storage Service

QUESTION 12

Ansible supports running Playbook on the host directly or via SSH. How can Ansible be told to run its playbooks directly on the host?

A. Setting `connection: local\\' in the tasks that run locally.

B. Specifying `-type local\\' on the command line.

C. It does not need to be specified; it is the default.

D. Setting `connection: local\\' in the Playbook.

Correct Answer: D

Ansible can be told to run locally on the command line with the `-c\\' option or can be told via the `connection: local\\' declaration in the playbook. The default connection method is `remote\\'.

Reference: http://docs.ansible.com/ansible/intro_inventory.html#non-ssh-connection-types

QUESTION 13

A company has 20 service learns Each service team is responsible for its own microservice. Each service team uses a separate AWS account for its microservice and a VPC with the 192 168 0 0/22 CIDR block. The company manages the AWS accounts with AWS Organizations.

Each service team hosts its microservice on multiple Amazon EC2 instances behind an Application Load Balancer. The microservices communicate with each other across the public internet. The company\\'s security team has issued a new guideline that all communication between microservices must use HTTPS over private network connections and cannot traverse the public internet.

A DevOps engineer must implement a solution that fulfills these obligations and minimizes the number of changes for each service team

Which solution will meet these requirements?

A. Create a new AWS account in AWS Organizations Create a VPC in this account and use AWS Resource Access Manager to share the private subnets of this VPC with the organization Instruct the service teams to launch a new. Network Load Balancer (NLB) and EC2 instances that use the shared private subnets Use the NLB DNS names for communication between microservices.

B. Create a Network Load Balancer (NLB) in each of the microservice VPCs Use AWS PrivateLink to create VPC endpoints in each AWS account for the NLBs Create subscriptions to each VPC endpoint in each of the other AWS accounts Use the VPC endpoint DNS names for communication between microservices.

C. Create a Network Load Balancer (NLB) in each of the microservice VPCs Create VPC peering connections between each of the microservice VPCs Update the route tables for each VPC to use the peering links Use the NLB DNS names for communication between microservices.

D. Create a new AWS account in AWS Organizations Create a transit gateway in this account and use AWS Resource Access Manager to share the transit gateway with the organization. In each of the microservice VPCs. create a transit gateway attachment to the shared transit gateway Update the route tables of each VPC to use the transit gateway Create a Network Load Balancer (NLB) in each of the microservice VPCs Use the NLB DNS names for communication between microservices.

Correct Answer: B

https://aws.amazon.com/blogs/networking-and-content-delivery/connecting-networks-with-overlapping-ip-ranges/
Private link is the best option because Transit Gateway doesn\\'t support overlapping CIDR ranges.

QUESTION 14

You have an application which consists of EC2 instances in an Auto Scaling group. Between a particular time frame every day, there is an increase in traffic to your website. Hence users are complaining of a poor response time on the

application. You have configured your Auto Scaling group to deploy one new EC2 instance when CPU utilization is greater than 60% for 2 consecutive periods of 5 minutes.

What is the least cost-effective way to resolve this problem?

A. Decrease the consecutive number of collection periods

B. Increase the minimum number of instances in the Auto Scaling group

C. Decrease the collection period to ten minutes

D. Decrease the threshold CPU utilization percentage at which to deploy a new instance

Correct Answer: B

If you increase the minimum number of instances, then they will be running even though the load is not high on the website. Hence you are incurring cost even though there is no need. All of the remaining options are possible options which can be used to increase the number of instances on a high load. For more information on On-demand scaling, please refer to the below link.

Reference: http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scale-based-on-demand.html

QUESTION 15

A company manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The EC2 instances run in an Auto Scaling group across multiple Availability Zones. The application uses an Amazon RDS for MySQL DB instance to store the data. The company has configured Amazon Route 53 with an alias record that points to the ALB.

A new company guideline requires a geographically isolated disaster recovery (DR> site with an RTO of 4 hours and an RPO of 15 minutes. Which DR strategy will meet these requirements with the LEAST change to the application stack?

A. Launch a replica environment of everything except Amazon RDS in a different Availability Zone Create an RDS read replica in the new Availability Zone: and configure the new stack to point to the local RDS DB instance. Add the new stack to the Route 53 record set by using a hearth check to configure a failover routing policy.

B. Launch a replica environment of everything except Amazon RDS in a different AWS. Region Create an RDS read replica in the new Region and configure the new stack to point to the local RDS DB instance. Add the new stack to the Route 53 record set by using a health check to configure a latency routing policy.

C. Launch a replica environment of everything except Amazon RDS ma different AWS Region. In the event of an outage copy and restore the latest RDS snapshot from the primary. Region to the DR Region Adjust the Route 53 record set to point to the ALB in the DR Region.

D. Launch a replica environment of everything except Amazon RDS in a different AWS Region. Create an RDS read replica in the new Region and configure the new environment to point to the local RDS DB instance. Add the new stack to the Route 53 record set by using a health check to configure a failover routing policy. In the event of an outage promote the read replica to primary.

Correct Answer: D