

KillTest

高品質 より良いサービス



練習問題

<https://www.killtest.jp>

1年間無料アップデート

Exam : SAP-C01

**Title : AWS Certified Solutions
Architect – Professional**

Version : DEMO

1. A company has a three-tier application running on AWS with a web server, an application server, and an Amazon RDS MySQL DB instance. A solutions architect is designing a disaster recovery (DR) solution with an RPO of 5 minutes.

Which solution will meet the company's requirements?

- A. Configure AWS Backup to perform cross-Region backups of all servers every 5 minutes. Reprovision the three tiers in the DR Region from the backups using AWS CloudFormation in the event of a disaster.
- B. Maintain another running copy of the web and application server stack in the DR Region using AWS CloudFormation drift detection. Configure cross-Region snapshots of the DB instance to the DR Region every 5 minutes. In the event of a disaster, restore the DB instance using the snapshot in the DR Region.
- C. Use Amazon EC2 Image Builder to create and copy AMIs of the web and application server to both the primary and DR Regions. Create a cross-Region read replica of the DB instance in the DR Region. In the event of a disaster, promote the read replica to become the master and reprovision the servers with AWS CloudFormation using the AMIs.
- D. Create AMIs of the web and application servers in the DR Region. Use scheduled AWS Glue jobs to synchronize the DB instance with another DB instance in the DR Region. In the event of a disaster, switch to the DB instance in the DR Region and reprovision the servers with AWS CloudFormation using the AMIs.

Answer: C

Explanation:

deploying a brand new RDS instance will take >30 minutes. You will use EC2 Image builder to put the AMIs into the new region, but not use image builder to LAUNCH them.

2. A solutions architect is building a web application that uses an Amazon RDS for PostgreSQL DB instance. The DB instance is expected to receive many more reads than writes. The solutions architect needs to ensure that the large amount of read traffic can be accommodated and that the DB instance is highly available.

Which steps should the solutions architect take to meet these requirements? (Select THREE.)

- A. Create multiple read replicas and put them into an Auto Scaling group
- B. Create multiple read replicas in different Availability Zones.
- C. Create an Amazon Route 53 hosted zone and a record set for each read replica with a TTL and a weighted routing policy
- D. Create an Application Load Balancer (ALB) and put the read replicas behind the ALB.
- E. Configure an Amazon CloudWatch alarm to detect a failed read replica. Set the alarm to directly invoke an AWS Lambda function to delete its Route 53 record set.
- F. Configure an Amazon Route 53 health check for each read replica using its endpoint

Answer: B,C,F

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/requests-rds-read-replicas/> You can use Amazon Route 53 weighted record sets to distribute requests across your read replicas. Within a Route 53 hosted zone, create individual record sets for each DNS endpoint associated with your read replicas and give them the same weight. Then, direct requests to the endpoint of the record set. You can incorporate Route 53 health checks to be sure that Route 53 directs traffic away from unavailable read replicas

3.A company has a new application that needs to run on five Amazon EC2 instances in a single AWS Region. The application requires high-throughput, low-latency network connections between all of the EC2 instances where the application will run. There is no requirement for the application to be fault tolerant.

Which solution will meet these requirements?

- A. Launch five new EC2 instances into a cluster placement group. Ensure that the EC2 instance type supports enhanced networking.
- B. Launch five new EC2 instances into an Auto Scaling group in the same Availability Zone. Attach an extra elastic network interface to each EC2 instance.
- C. Launch five new EC2 instances into a partition placement group. Ensure that the EC2 instance type supports enhanced networking.
- D. Launch five new EC2 instances into a spread placement group. Attach an extra elastic network interface to each EC2 instance.

Answer: A

Explanation:

When you launch EC2 instances in a cluster they benefit from performance and low latency. No redundancy though as per the question

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

4.A company needs to implement a patching process for its servers. The on-premises servers and Amazon EC2 instances use a variety of tools to perform patching. Management requires a single report showing the patch status of all the servers and instances.

Which set of actions should a solutions architect take to meet these requirements?

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

Answer: A

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-patch.html>

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

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

- A. Store a random string in AWS Secrets Manager. Create an AWS Lambda function for automatic secret rotation. Configure CloudFront to inject the random string as a custom HTTP header for the origin

request. Create an AWS WAF web ACL rule with a string match rule for the custom header. Associate the web ACL with the ALB.

B. Create an AWS WAF web ACL rule with an IP match condition of the CloudFront service IP address ranges. Associate the web ACL with the ALB. Move the ALB into the three private subnets.

C. Store a random string in AWS Systems Manager Parameter Store. Configure Parameter Store automatic rotation for the string. Configure CloudFront to inject the random string as a custom HTTP header for the origin request. Inspect the value of the custom HTTP header, and block access in the ALB.

D. Configure AWS Shield Advanced. Create a security group policy to allow connections from CloudFront service IP address ranges. Add the policy to AWS Shield Advanced, and attach the policy to the ALB.

Answer: D

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html>

it shows For Amazon EC2 Auto Scaling, there are two primary process types: Launch and Terminate.

The Launch process adds a new Amazon EC2 instance to an Auto Scaling group, increasing its capacity.

The Terminate process removes an Amazon EC2 instance from the group, decreasing its capacity.

HealthCheck process for EC2 autoscaling is not a primary process! It is a process along with the

following AddToLoadBalancer AlarmNotification AZRebalance HealthCheck InstanceRefresh

ReplaceUnhealthy ScheduledActions. From the requirements, Some EC2 instances are now being

marked as unhealthy and are being terminated. Application is running at reduced capacity not because instances are marked unhealthy but because they are being terminated.

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html#choosing-suspend-resume>