

1) A media company is running a critical production application that uses Amazon RDS for PostgreSQL with Multi-AZ deployments. The database size is currently 25 TB. The IT director wants to migrate the database to Amazon Aurora PostgreSQL with minimal effort and minimal disruption to the business.

What is the best migration strategy to meet these requirements?

- A) Use the AWS Schema Conversion Tool (AWS SCT) to copy the database schema from RDS for PostgreSQL to an Aurora PostgreSQL DB cluster. Create an AWS DMS task to copy the data.
- B) Create a script to continuously back up the RDS for PostgreSQL instance using `pg_dump`, and restore the backup to an Aurora PostgreSQL DB cluster using `pg_restore`.
- C) Create a read replica from the existing production RDS for PostgreSQL instance. Check that the replication lag is zero and then promote the read replica as a standalone Aurora PostgreSQL DB cluster.
- D) Create an Aurora Replica from the existing production RDS for PostgreSQL instance. Stop the writes on the master, check that the replication lag is zero, and then promote the Aurora Replica as a standalone Aurora PostgreSQL DB cluster.

2) A medical company is planning to migrate its on-premises PostgreSQL database, along with application and web servers, to AWS. Amazon RDS for PostgreSQL is being considered as the target database engine. Access to the database should be limited to application servers and a bastion host in a VPC.

Which solution meets the security requirements?

- A) Launch the RDS for PostgreSQL database in a DB subnet group containing private subnets. Modify the `pg_hba.conf` file on the DB instance to allow connections from only the application servers and bastion host.
- B) Launch the RDS for PostgreSQL database in a DB subnet group containing public subnets. Create a new security group with inbound rules to allow connections from only the security groups of the application servers and bastion host. Attach the new security group to the DB instance.
- C) Launch the RDS for PostgreSQL database in a DB subnet group containing private subnets. Create a new security group with inbound rules to allow connections from only the security groups of the application servers and bastion host. Attach the new security group to the DB instance.
- D) Launch the RDS for PostgreSQL database in a DB subnet group containing private subnets. Create a NACL attached to the VPC and private subnets. Modify the inbound and outbound rules to allow connections to and from the application servers and bastion host.

3) A database specialist is troubleshooting complaints from an application's users who are experiencing performance issues when saving data in an Amazon ElastiCache for Redis cluster with cluster mode disabled. The database specialist finds that the performance issues are occurring during the cluster's backup window. The cluster runs in a replication group containing three nodes. Memory on the nodes is fully utilized. Organizational policies prohibit the database specialist from changing the backup window time.

How could the database specialist address the performance concern? (Select TWO.)

- A) Add an additional node to the cluster in the same Availability Zone as the primary.
- B) Configure the backup job to take a snapshot of a read replica.
- C) Increase the local instance storage size for the cluster nodes.
- D) Increase the reserved-memory-percent parameter value.
- E) Configure the backup process to flush the cache before taking the backup.

4) A company's security department has mandated that their existing Amazon RDS for MySQL DB instance be encrypted at rest.

What should a database specialist do to meet this requirement?

- A) Modify the database to enable encryption. Apply this setting immediately without waiting for the next scheduled maintenance window.
- B) Export the database to an Amazon S3 bucket with encryption enabled. Create a new database and import the export file.
- C) Create a snapshot of the database. Create an encrypted copy of the snapshot. Create a new database from the encrypted snapshot.
- D) Create a snapshot of the database. Restore the snapshot into a new database with encryption enabled.

5) A company has a highly available production 10 TB SQL Server relational database running on Amazon EC2. Users have recently been reporting performance and connectivity issues. A database specialist has been asked to configure a monitoring and alerting strategy that will provide metrics visibility and notifications to troubleshoot these issues.

Which solution will meet these requirements?

- A) Configure AWS CloudTrail logs to monitor and detect signs of potential problems. Create an AWS Lambda function that is triggered when specific API calls are made and send notifications to an Amazon SNS topic.
- B) Install an Amazon Inspector agent on the DB instance. Configure the agent to stream server and database activity to Amazon CloudWatch Logs. Configure metric filters and alarms to send notifications to an Amazon SNS topic.
- C) Migrate the database to Amazon RDS for SQL Server and use Performance Insights to monitor and detect signs of potential problems. Create a scheduled AWS Lambda function that retrieves metrics from the Performance Insights API and send notifications to an Amazon SNS topic.
- D) Configure Amazon CloudWatch Application Insights for .NET and SQL Server to monitor and detect signs of potential problems. Configure CloudWatch Events to send notifications to an Amazon SNS topic.

6) A company's ecommerce application stores order transactions in an Amazon RDS for MySQL database. The database has run out of available storage and the application is currently unable to take orders.

Which action should a database specialist take to resolve the issue in the shortest amount of time?

- A) Add more storage space to the DB instance using the ModifyDBInstance action.
- B) Create a new DB instance with more storage space from the latest backup.
- C) Change the DB instance status from STORAGE_FULL to AVAILABLE.
- D) Configure a read replica with more storage space.

7) A company undergoing a security audit has determined that its database administrators are presently sharing an administrative database user account for the company's Amazon Aurora deployment. To support proper traceability, governance, and compliance, each database administration team member must start using individual, named accounts. Furthermore, long-term database user credentials should not be used.

Which solution should a database specialist implement to meet these requirements?

- A) Use the AWS CLI to fetch the AWS IAM users and passwords for all team members. For each IAM user, create an Aurora user with the same password as the IAM user.
- B) Enable IAM database authentication on the Aurora cluster. Create a database user for each team member without a password. Attach an IAM policy to each administrator's IAM user account that grants the connect privilege using their database user account.
- C) Create a database user for each team member. Share the new database user credentials with the team members. Have users change the password on the first login to the same password as their IAM user.
- D) Create an IAM role and associate an IAM policy that grants the connect privilege using the shared account. Configure a trust policy that allows the administrator's IAM user account to assume the role.

8) A global company wants to run an application in several AWS Regions to support a global user base. The application will need a database that can support a high volume of low-latency reads and writes that is expected to vary over time. The data must be shared across all of the Regions to support dynamic company-wide reports.

Which database meets these requirements?

- A) Use Amazon Aurora Serverless and configure endpoints in each Region.
- B) Use Amazon RDS for MySQL and deploy read replicas in an auto scaling group in each Region.
- C) Use Amazon DocumentDB (with MongoDB compatibility) and configure read replicas in an auto scaling group in each Region.
- D) Use Amazon DynamoDB global tables and configure DynamoDB auto scaling for the tables.

9) A company's customer relationship management application uses an Amazon RDS for PostgreSQL Multi-AZ database. The database size is approximately 100 GB. A database specialist has been tasked with developing a cost-effective disaster recovery plan that will restore the database in a different Region within 2 hours. The restored database should not be missing more than 8 hours of transactions.

What is the MOST cost-effective solution that meets the availability requirements?

- A) Create an RDS read replica in the second Region. For disaster recovery, promote the read replica to a standalone instance.
- B) Create an RDS read replica in the second Region using a smaller instance size. For disaster recovery, scale the read replica and promote it to a standalone instance.
- C) Schedule an AWS Lambda function to create an hourly snapshot of the DB instance and another Lambda function to copy the snapshot to the second Region. For disaster recovery, create a new RDS Multi-AZ DB instance from the last snapshot.
- D) Create a new RDS Multi-AZ DB instance in the second Region. Configure an AWS DMS task for ongoing replication.

10) An operations team in a large company wants to centrally manage resource provisioning for its development teams across multiple accounts. When a new AWS account is created, the developers require full privileges for a database environment that uses the same configuration, data schema, and source data as the company's production Amazon RDS for MySQL DB instance.

How can the operations team achieve this?

- A) Enable the source DB instance to be shared with the new account so the development team may take a snapshot. Create an AWS CloudFormation template to launch the new DB instance from the snapshot.
- B) Create an AWS CLI script to launch the approved DB instance configuration in the new account. Create an AWS DMS task to copy the data from the source DB instance to the new DB instance.
- C) Take a manual snapshot of the source DB instance and share the snapshot privately with the new account. Specify the snapshot ARN in an RDS resource in an AWS CloudFormation template and use StackSets to deploy to the new account.
- D) Create a DB instance read replica of the source DB instance. Share the read replica with the new AWS account.

Answers

- 1) D - To migrate from an Amazon RDS for PostgreSQL DB instance to an [Amazon Aurora PostgreSQL DB cluster](#), create an Aurora Replica of your source PostgreSQL DB instance. When the replica lag between the PostgreSQL DB instance and the Aurora PostgreSQL Replica is zero, you can promote the Aurora Replica to be a standalone Aurora PostgreSQL DB cluster.
- 2) C - [Create custom rules](#) in the security group for your DB instances that allow connections from the security group you created for your Amazon EC2 instances. This would allow instances associated with the security group to access the DB instances. Including [bastion hosts](#) in your VPC environment enables you to securely connect to your database instances running in private subnets.
- 3) B,D - Because of the system resources required during a backup, we recommend that you [create backups from one of the read replicas](#). If all of a node's available memory is consumed, then excessive paging to the disk can occur. We recommend setting the reserved-memory-percent [parameter to 25%](#) to reserve enough memory for background processes.
- 4) C - To [enable encryption at rest](#) for an existing unencrypted DB instance, you can create a snapshot of your DB instance, and then create an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot so you have an encrypted copy of your original DB instance.
- 5) D - [Amazon CloudWatch Application Insights](#) uses machine learning classification algorithms to analyze metrics and identify signs of problems with your applications. Windows Event Viewer and SQL Server Error logs are included in the analysis. To receive notifications, you can create an Amazon EventBridge (in CloudWatch Events) [rule for the Application Insights Problem Detected event](#).
- 6) A - If your DB instance runs out of storage space, it might no longer be available. To recover from this scenario, [add more storage space](#) to your instance using the ModifyDBInstance action. To prevent storage space issues from happening in the future, [enable storage autoscaling](#).
- 7) B - You can authenticate to your DB cluster using [AWS IAM database authentication](#). With this authentication method, you don't need to use a password when you connect to a DB cluster. Instead, you use an authentication token that expires 15 minutes after creation.
- 8) D - [Amazon DynamoDB global tables](#) provide a multi-Region, multi-master database in the AWS Regions you specify. DynamoDB performs all of the necessary tasks to create identical tables in these Regions and propagate ongoing data changes to all of them. [DynamoDB auto scaling](#) cost effectively adjusts provisioned throughput to actual traffic patterns.
- 9) C - Backup and restore is the most [cost-effective solution](#) to provide a 2-hour RTO and 8-hour RPO. Manual hourly snapshots need to be copied to the second Region to be available for the creation of the new database. Taking the snapshots every hour will keep the incremental snapshot size low, reduce the time to copy the snapshot across Regions, and meet the RPO. Also, taking snapshots frequently does not impact the cost. A [pair of AWS Lambda functions](#) can be scheduled to take the snapshot and copy it to the second Region.
- 10) C - A manual DB snapshot [can be shared privately](#) with other AWS accounts. [AWS CloudFormation StackSets](#) extends the functionality of stacks by enabling you to create, update, or delete stacks across multiple accounts and Regions with a single operation. Using an administrator account, you define and manage an AWS CloudFormation template, and use the template as the basis for provisioning stacks into selected target accounts across specified Regions.