Guidance for Disaster Recovery Using Amazon Aurora

This architecture diagram shows how to implement an Aurora Global database to replicate data to a secondary Region.

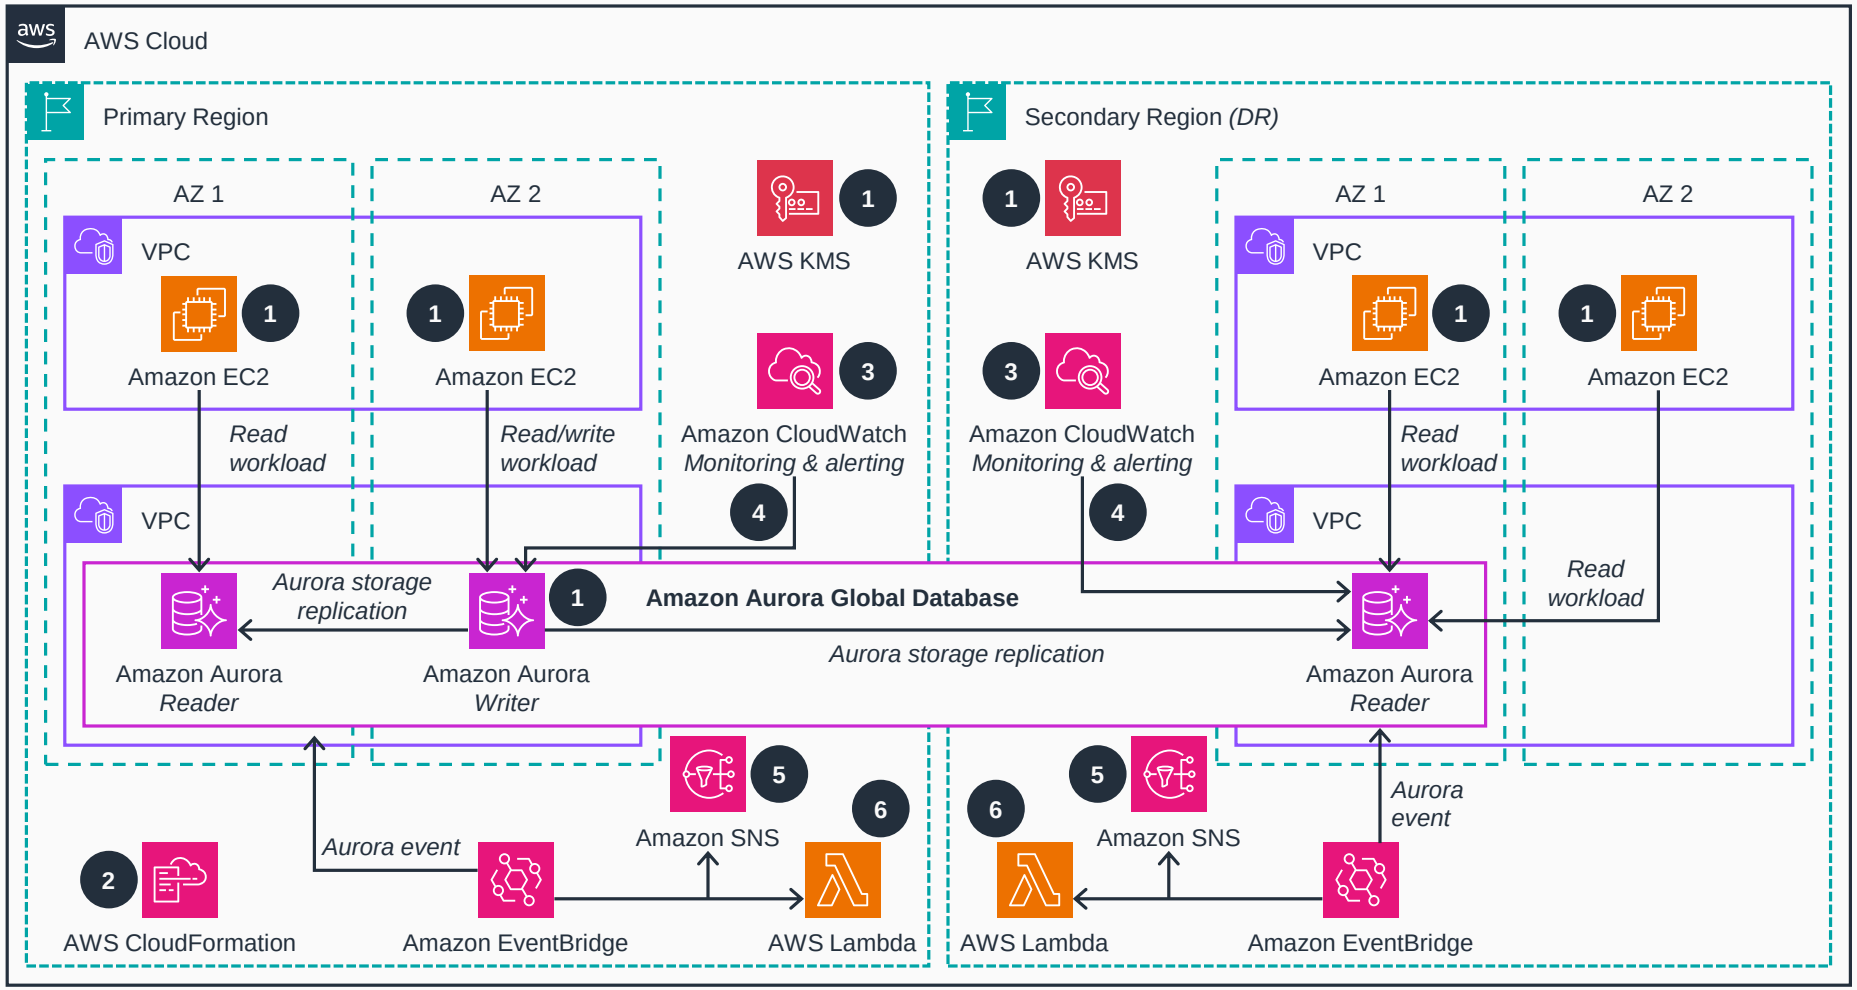

1. **Prerequisites:** This Guidance requires an existing **Amazon Aurora** Regional cluster. The application can run on **Amazon Elastic Compute Cloud (Amazon EC2)**, **Amazon Elastic Kubernetes Service (Amazon EKS)**, **Amazon Elastic Container Service (Amazon ECS)**, or another service of your choice. This Guidance assumes you have used **Amazon EC2** instances in virtual private clouds (VPCs) across multiple Availability Zones (AZs). You can encrypt an **Aurora** cluster using the default **AWS Key Management Service (AWS KMS)** or using a customer-managed key (CMK).

2. **AWS CloudFormation** creates resources, including an **Aurora** read replica in the primary AWS Region if one does not exist already and an **Aurora** global database with a reader instance in the secondary Region. An **Amazon CloudWatch** dashboard, an **Amazon Simple Notification Service (Amazon SNS)** topic, an **AWS Lambda** function, and **Amazon EventBridge** rules are deployed in both Regions.

3. The **CloudWatch** dashboard is configured in the primary and secondary Regions to monitor key metrics related to **Aurora**, along with the replication status.

4. A **CloudWatch** alarm is created in both Regions to generate alarms for AuroraGlobalDBReplicationLag metrics and notifications through the **Amazon SNS** topic.

5. An **EventBridge** rule is configured for planned switchovers and unplanned failovers. When an event occurs, it sends notifications using **Amazon SNS** and calls the **Lambda** functions in both Regions.

6. The **Lambda** function provides a framework to add any additional functionalities during the failover event. For example:

   - You can configure the application to use an **Amazon Route 53** in a newly promoted Region during a failover event so that no application configuration is required during the event.
   - You can configure the application to restart an **Amazon EC2** instance or the application deployment pods in the **Amazon EKS** cluster after the database failover event.

Reviewed for technical accuracy February 19, 2024

AWS Reference Architecture

Guidance for Disaster Recovery Using Amazon Aurora

This architecture diagram shows how to backup your Amazon Aurora database automatically to another Region and/or to another account using AWS Backup. In the event of a disaster, these backups can be restored to create a new Aurora database.

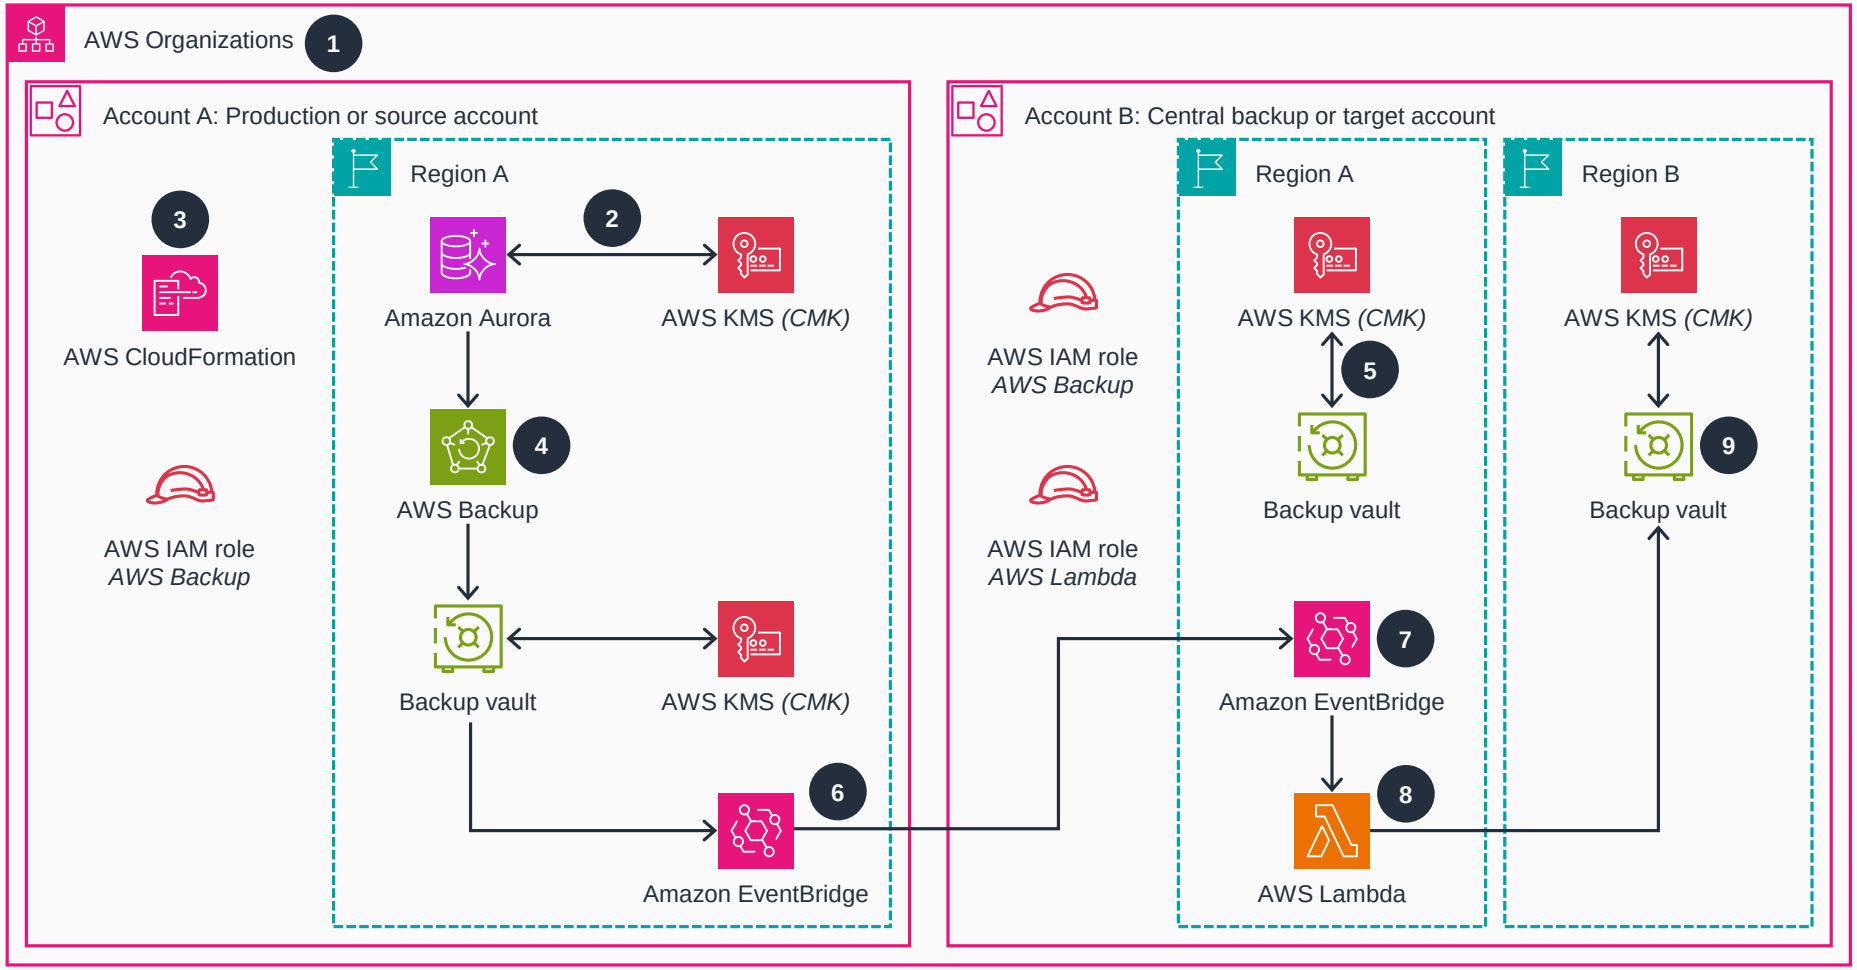

1. A preexisting organizational structure within **AWS Organizations** is necessary to establish cross-account **AWS Backup** between two accounts: Account A, which serves as the production or "source" account, and Account B, which is the central backup or "target" account. Notably, this Guidance provides the flexibility to include multiple target accounts.
2. An existing **Aurora** cluster in the source account is encrypted using a CMK that is shared across the source and target accounts. The cluster should also be tagged appropriately so that the solution can identify the desired resources for backup.
3. **CloudFormation** is used to deploy the solution resources in your source and target AWS accounts and Regions. The required **CloudFormation** stacks are provided as part of this solution.
4. The **Aurora** cluster in Region A of the production account is backed by **AWS Backup** according to the schedule you provided while deploying the solution. The backups are stored in an **AWS Backup** vault encrypted with an **AWS KMS** CMK.
5. **AWS Backup** copies the backups to the cross-account (that is, the target account) and stores it in the backup vault in Account B Region A. The backup vault is encrypted using a CMK created by **CloudFormation**.
6. Once the cross-account backup copy is complete, an **EventBridge** rule in the source account forwards a "backup copy complete" notification to the target account event bus (Account B Region A).
7. An **EventBridge** rule in the target account in Region A identifies the notification as an incoming event.
8. Once the event is received in the target account, the **EventBridge** rule invokes a **Lambda** function to finally copy the backup to the desired destination (Account B Region B) and store it in the **AWS Backup** vault in Region B.
9. The backup of your **Aurora** cluster is now available in the target account in Region B and is stored in the **AWS Backup** vault of Region B. The backup vault is encrypted with an **AWS KMS** CMK. This backup can be used to restore the **Aurora** database.

Reviewed for technical accuracy February 19, 2024

AWS Reference Architecture