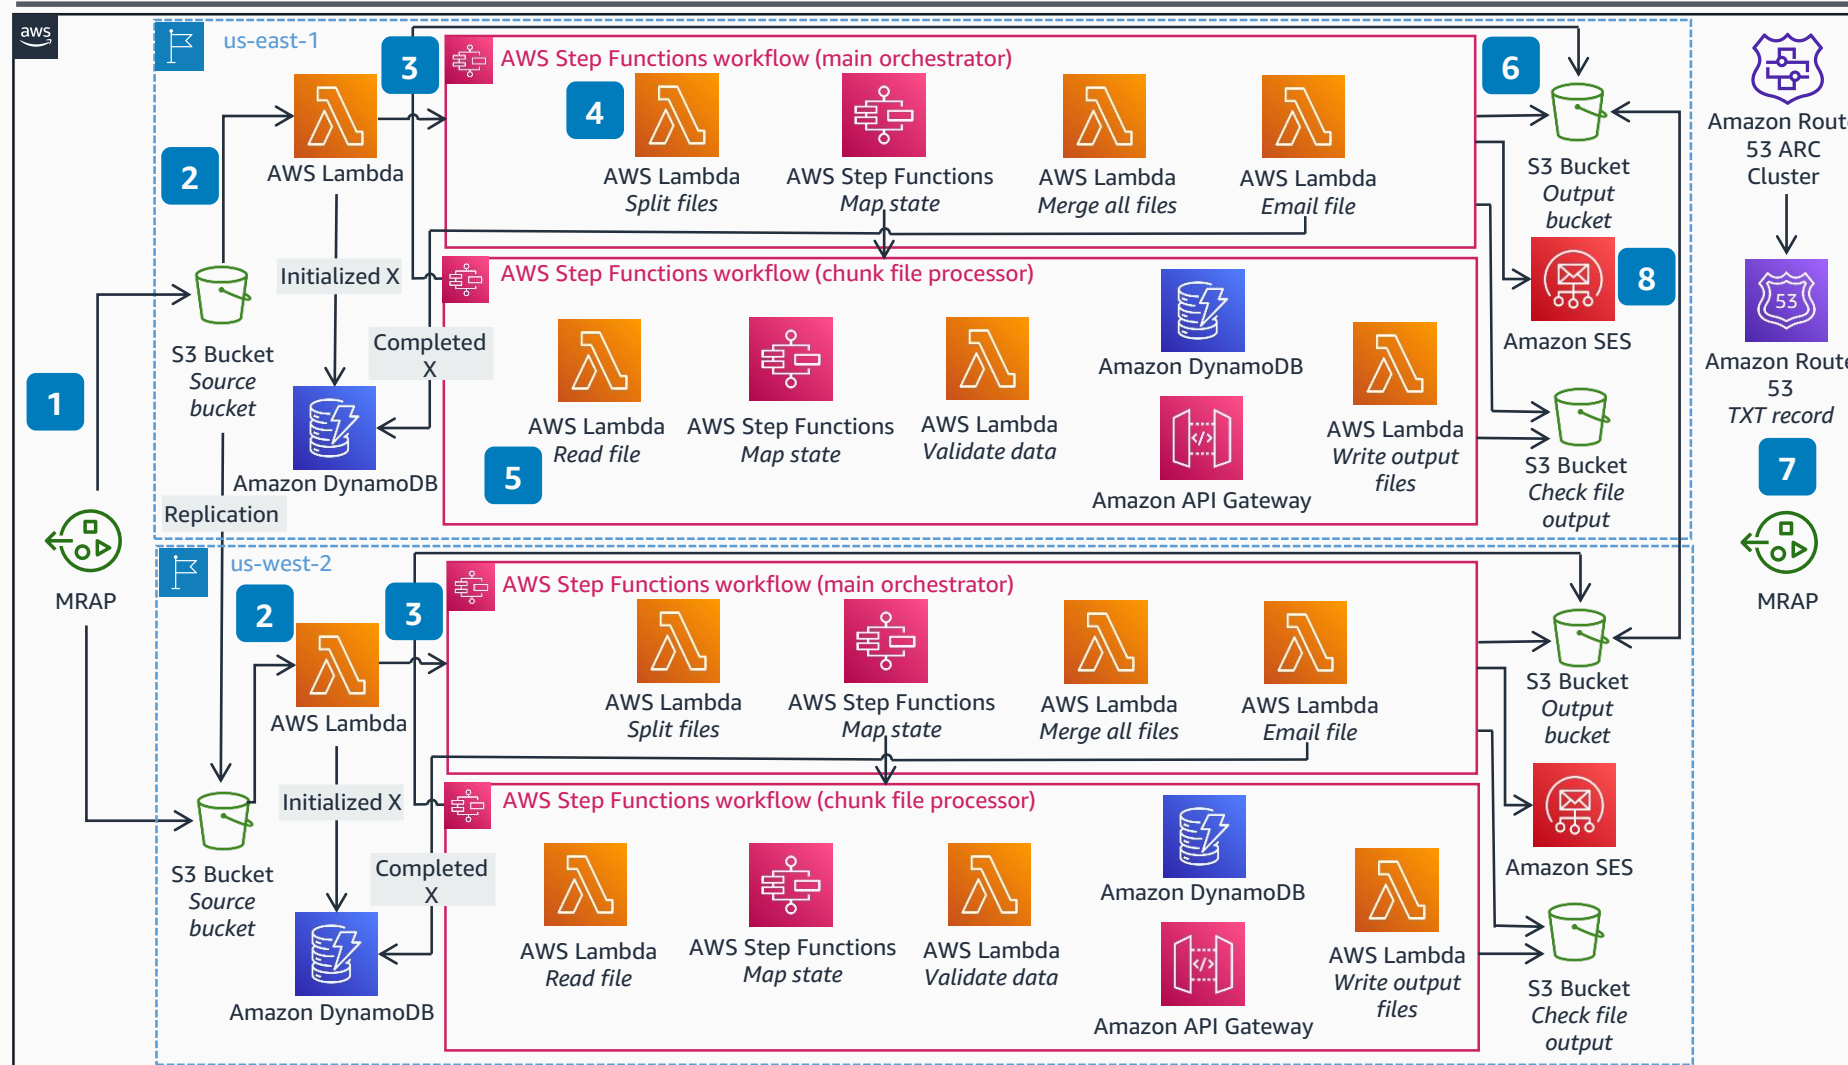


Guidance for Multi-Region Serverless Batch Application on AWS

This architecture shows the multi-Region, event-driven workload when running in the primary Region.



- 1** Add the file to an **Amazon Simple Storage Service (Amazon S3)** bucket using the MRAP. MRAP routes the file to one of the **S3** buckets, which will replicate the object to the other bucket.
- 2** **Amazon S3** invokes the **AWS Lambda** function `putObject` in both Regions.
- 3** The **Lambda** function will resolve the TXT record in an **Amazon Route53** private hosted zone to determine if it is the active Region. If it is, the workflow will continue. If it is not, the function will exit and not take further action. The function in the active Region writes metadata on the file to the **Amazon DynamoDB** batch state table, including that the processing has started, and starts the first **AWS Step Functions** workflow.
- 4** The main orchestrator **Step Functions** workflow orchestrates file processing by splitting the file into small chunks and then passing it to chunk file processor **Step Functions** workflow.
- 5** The chunk file processor **Step Functions** workflow is responsible for processing each row from the chunk file.
- 6** The merged file is written to **Amazon S3**, which replicates it to the standby Region's bucket.
- 7** A pre-signed URL is generated using the MRAP so the user can retrieve the file from the closest **S3** bucket. The routing logic is abstracted from the client.
- 8** **Amazon Simple Email Service (Amazon SES)** mails the pre-signed URL to recipients so they can retrieve the file from one of the **S3** buckets through the MRAP.

