Guidance for Oracle's JD Edwards EnterpriseOne with AWS Elastic Disaster Recovery

This architecture diagram shows how to use AWS Elastic Disaster Recovery (AWS DRS) for disaster recovery using Oracle's JD Edwards EnterpriseOne software.

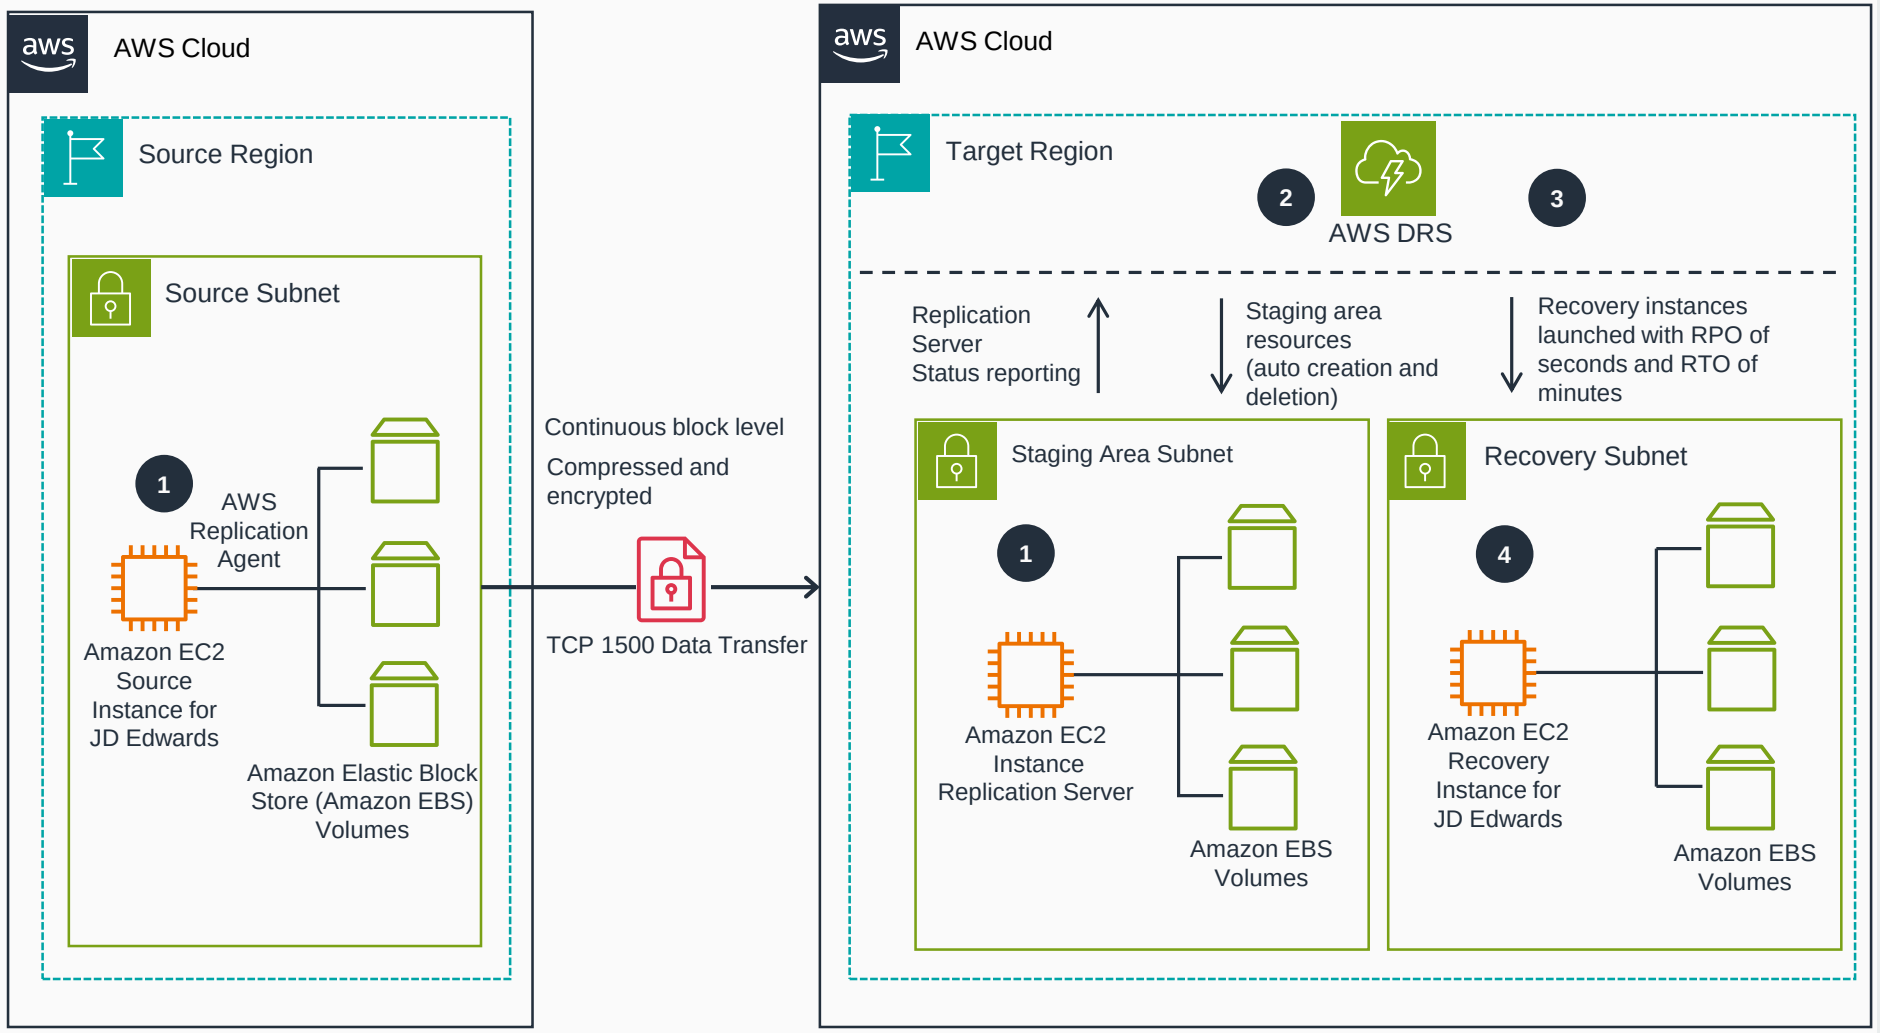

1. **AWS Elastic Disaster Recovery** replication begins with an initial sync. During the initial sync, the AWS Replication Agent replicates all the data from the source disks to the appropriate resource in the staging area subnet.

   Continuous replication continues indefinitely after the initial sync is complete.

2. You review the launch parameters, which include service-specific configurations and an **Amazon Elastic Compute Cloud** (Amazon EC2) launch template. When the source server is indicated as being ready for recovery, you can start instances.

   When **AWS DRS** issues a series of API calls to begin the launch operation, the recovery instance is immediately launched on AWS according to your launch settings.

3. The new instance is spun up on AWS after the conversion is complete and is ready for use. The conversion process involves changes to the drivers, network, and operating system license to ensure that the instance boots natively on AWS. The source server state at the time of launch is represented by the volumes associated with the launched instance.

4. After the launch, the newly created volumes are no longer kept in sync with the source servers. The AWS Replication Agent continues to routinely replicate changes made to your source servers to the staging area volumes, but the launched instances do not reflect those changes.

Reviewed for technical accuracy December 14, 2023

AWS Reference Architecture