Guidance for Scaling Electronic Design Automation (EDA) on AWS

This architecture diagram shows how to provide compute capacity with design data caching using a multi-Region approach.

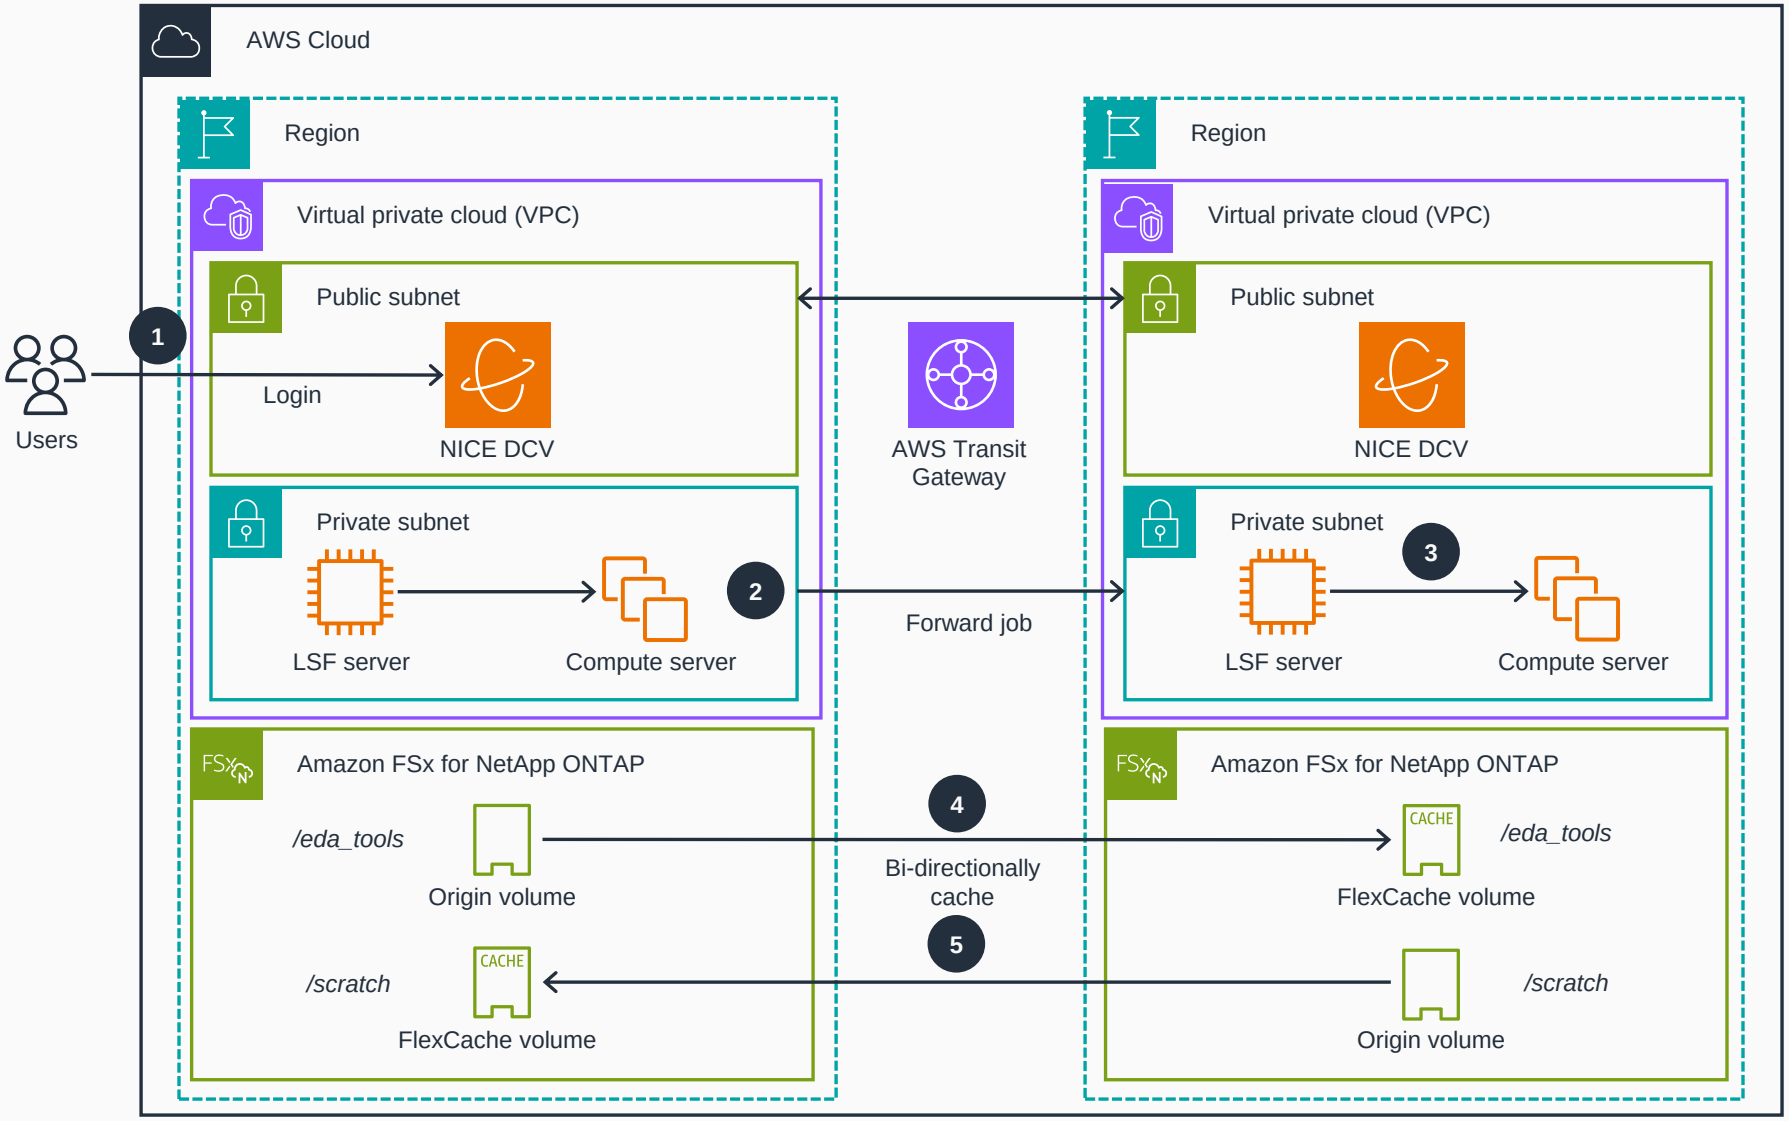

1. The user logs in to the **NICE DCV** desktop to submit the job through a multi-cluster of IBM Spectrum LSF, which is deployed using a multi-Region (or hybrid cloud) approach.
2. The peered LSF server forwards the job to the receiver queue in the other LSF primary when the user submits a job in the sender queue.
3. Once the job is received, the LSF server spins up the **Amazon Elastic Compute Cloud (Amazon EC2)** instance based on the resource requirements of the user job and the defined **Amazon EC2** resource configuration by using **EC2** fleet.
4. The user's job reads all related data from **Amazon FSx for NetApp ONTAP** that mounted the FlexCache volume on the spun-up **EC2** instance.
5. The user job writes data to the **FSx for ONTAP** volume that has cache configured to the FlexCache volume. The data result is located in the same path for output (/scratch), so the user doesn't need to log in to another cluster to view the data.

Reviewed for technical accuracy July 12, 2024

AWS Reference Architecture