Guidance for Writing High Transaction Workloads on Amazon S3 Express One Zone

This architecture diagram demonstrates how the unique features of Amazon S3 Express One Zone can decrease network latency and make writing higher throughput transactions scalable and fast.

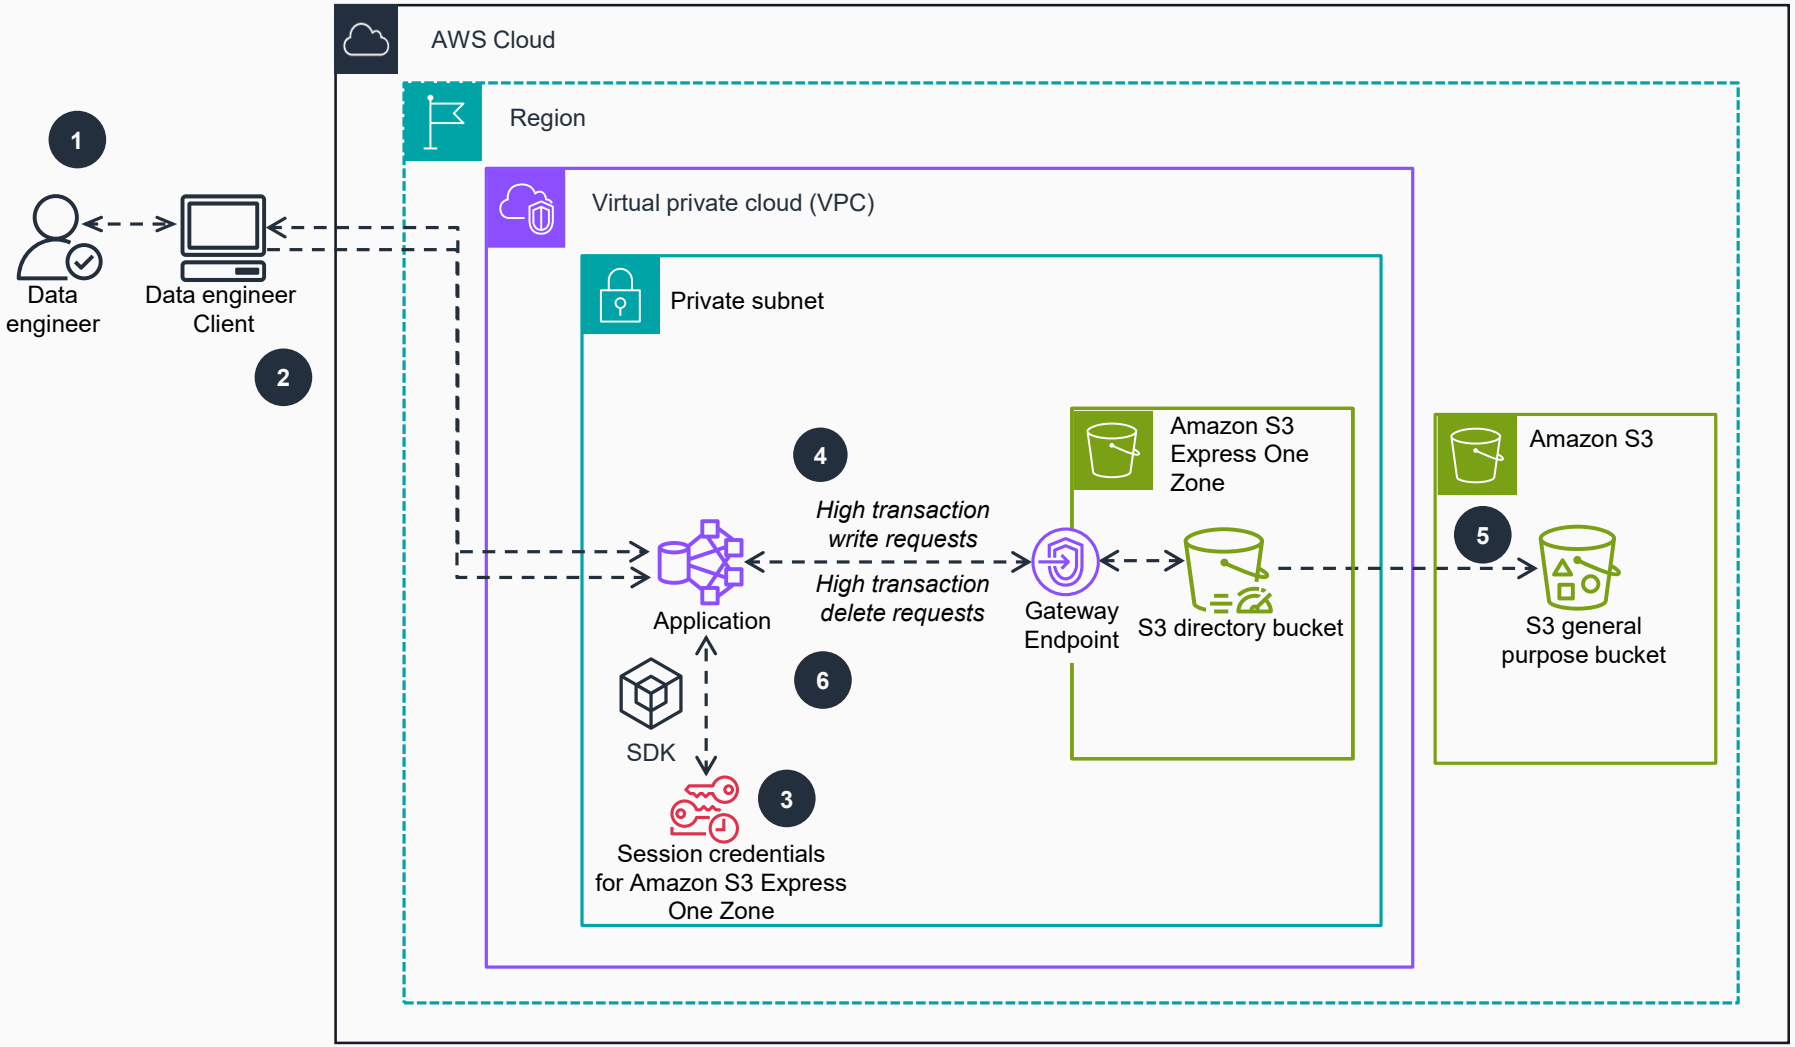

1. Data engineer creates an **Amazon Simple Storage Service (Amazon S3)** directory bucket for an application workload.
2. Data engineer creates or configures the appropriate policies for application access to the **Amazon S3 Express One Zone** storage class.
3. When the application interaction starts, session credentials are generated by an **S3 Express One Zone** Software Development Kit (SDK), authorized by **AWS Identity and Access Management** (IAM). The credentials can be used by an application for a limited time period of 5 minutes.
4. High-performance application data persistence *write* requests, done through **Amazon S3** API PUT requests, are directed to the **S3 Express One Zone** storage class through a gateway endpoint. Object data is available for high-performance query, extract, transform, and load (ETL), compaction, and more.
5. The application object data that needs to be persisted or no longer requires high-performance is copied to an appropriate **Amazon S3** storage class using an **Amazon S3** Batch Operations copy job.
6. You can delete data no longer needed from the **S3 Express One Zone** storage class through an authorized application or other options.

Reviewed for technical accuracy November 22, 2023

AWS Reference Architecture