Guidance for Getting Started with AWS Storage Tape Gateway

This architecture diagram shows you how to use your existing tape-based backup infrastructure to store data on virtual tape cartridges that you create on your Tape Gateway.

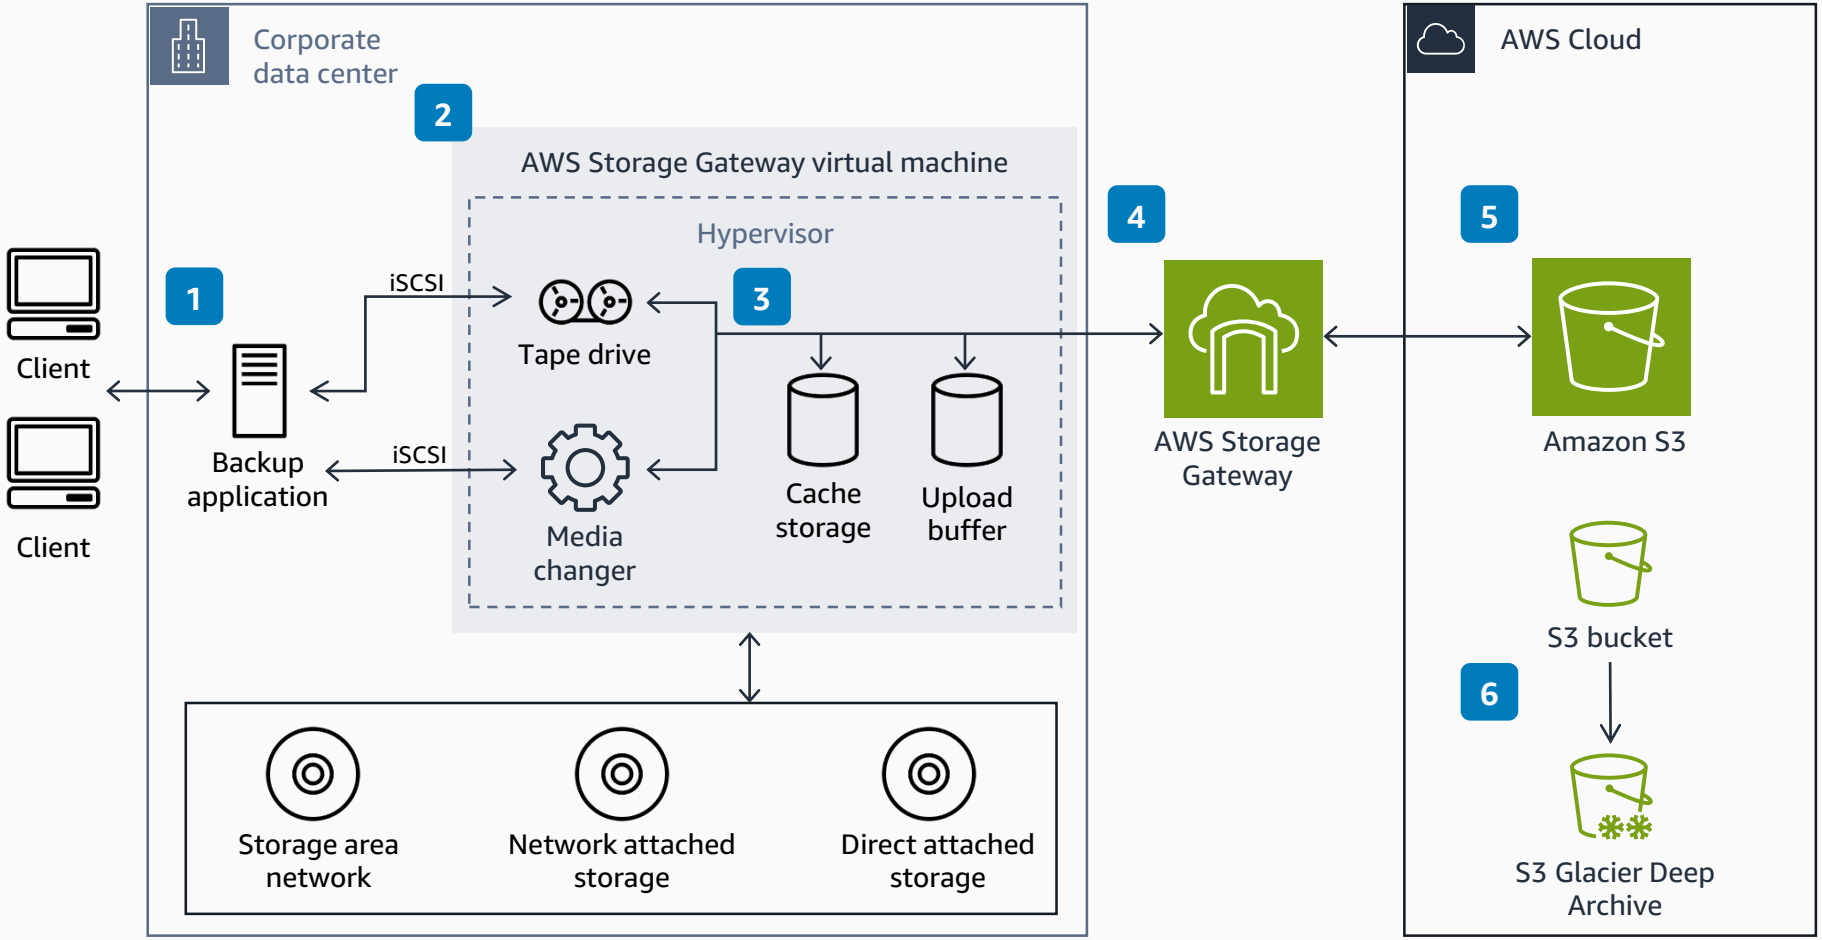

1. An on-premises backup application connects to your **AWS Storage Gateway** virtual machine through the Internet Small Computer Systems Interface (iSCSI).
2. A **Storage Gateway** virtual machine installed in your virtual environment acts as an interface to the AWS Cloud.
3. The cache storage acts as the durable store for any data awaiting upload from the upload buffer to **Amazon Simple Storage Service (Amazon S3)**. The upload buffer provides a staging area for the gateway before it uploads the data to a virtual tape.
4. The on-premises virtual machine connects to a **Storage Gateway** endpoint in your AWS account.
5. A virtual tape is similar to a physical tape cartridge. However, virtual tape data is stored in **Amazon S3**.
6. When your backup software ejects a tape, your gateway moves the tape to the archive for long-term storage. Tapes in the archive are stored in the virtual tape shelf (VTS). The VTS is backed by **S3 Glacier Flexible Retrieval** or **S3 Glacier Deep Archive**.

Reviewed for technical accuracy September 11, 2023

AWS Reference Architecture