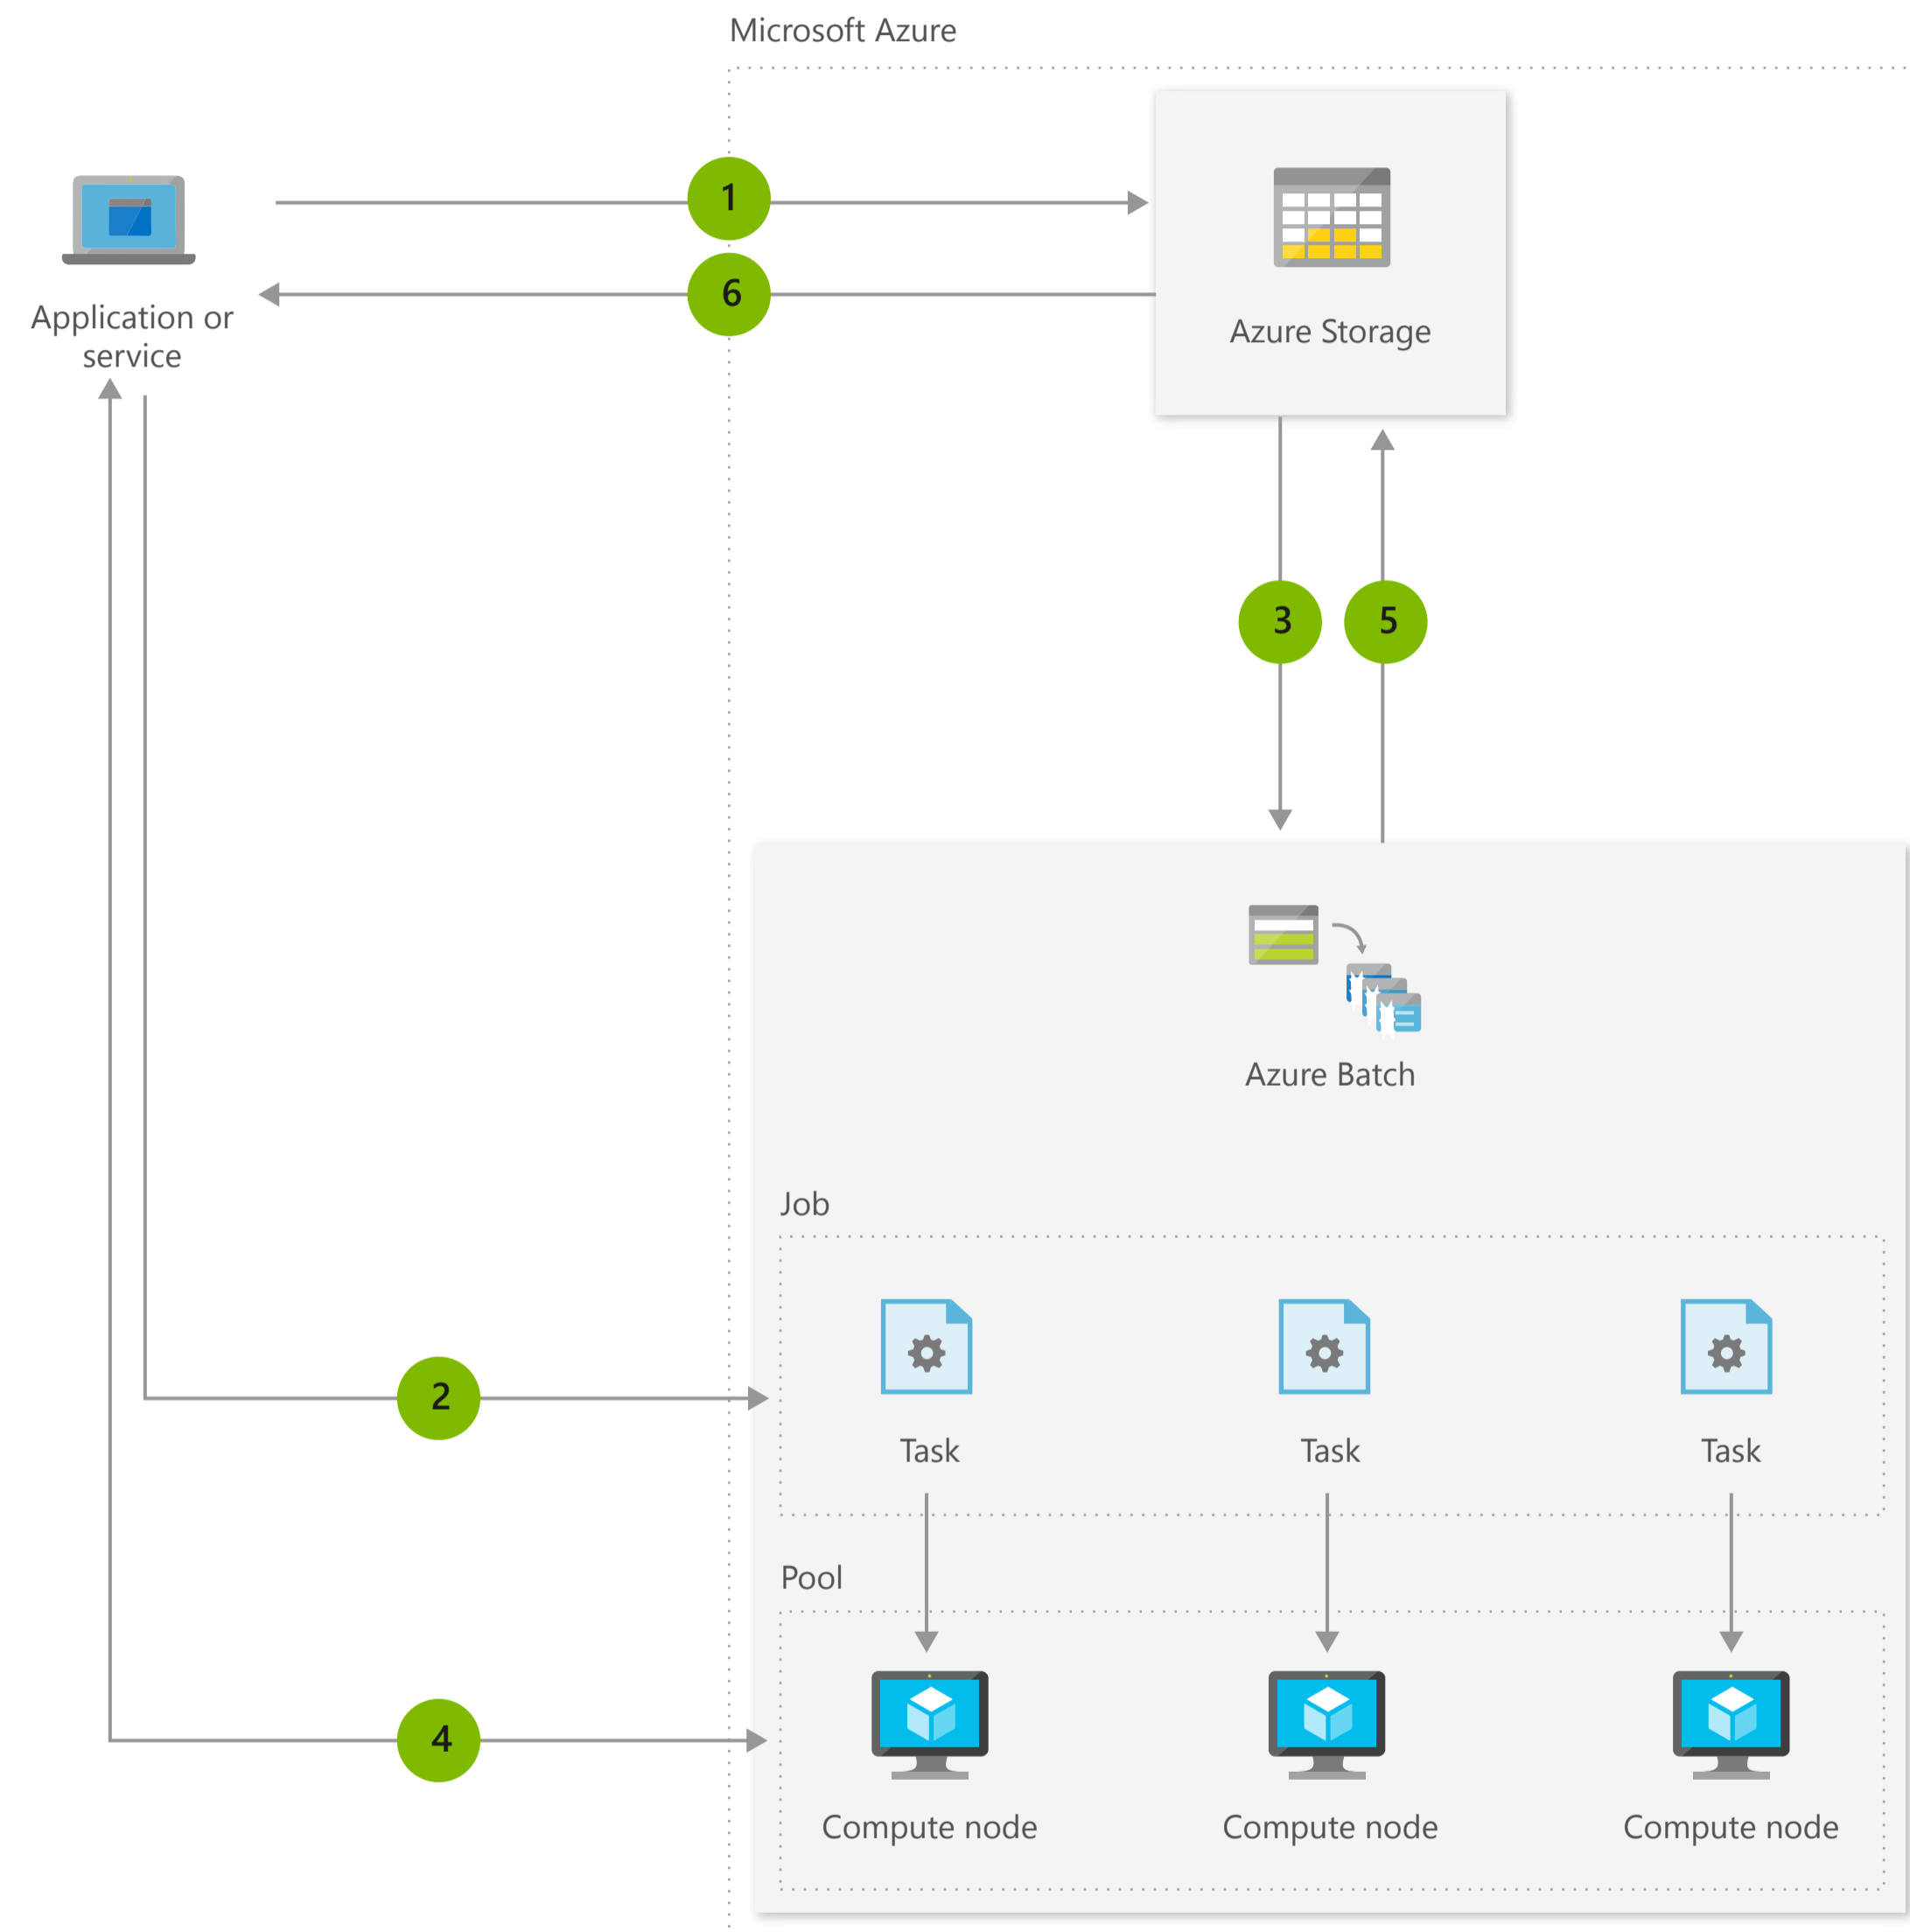


Big compute with Azure Batch



Architecture overview

Big compute and high performance computing (HPC) workloads are normally compute intensive and can be run in parallel, taking advantage of the scale and flexibility of the cloud. The workloads are often run asynchronously using batch processing, with compute resources required to run the work and job scheduling required to specify the work. Examples of Big Compute and HPC workloads include financial risk Monte Carlo simulations, image rendering, media transcoding, file processing, and engineering or scientific simulations.

This solution implements a cloud-native application with Azure Batch, which provides compute resource allocation and management, application installation, resource auto-scaling, and job scheduling as a platform service. Batch also offers higher level workload accelerators specifically for running R in parallel, AI training, and rendering workloads.

This solution is built on the Azure managed services—Virtual Machines, Storage, and Batch. These services run in a high-availability environment, patched and supported, allowing you to focus on your solution.

- 1 Upload input files and the applications to your Azure Storage account.
- 2 Create a Batch pool of compute nodes, a job to run the workload on the pool, and the tasks in the job.
- 3 Batch downloads input files and applications.
- 4 Batch monitors task execution.
- 5 Batch uploads task output.
- 6 Download output files.

Products

 Virtual Machines

 Batch

 Storage