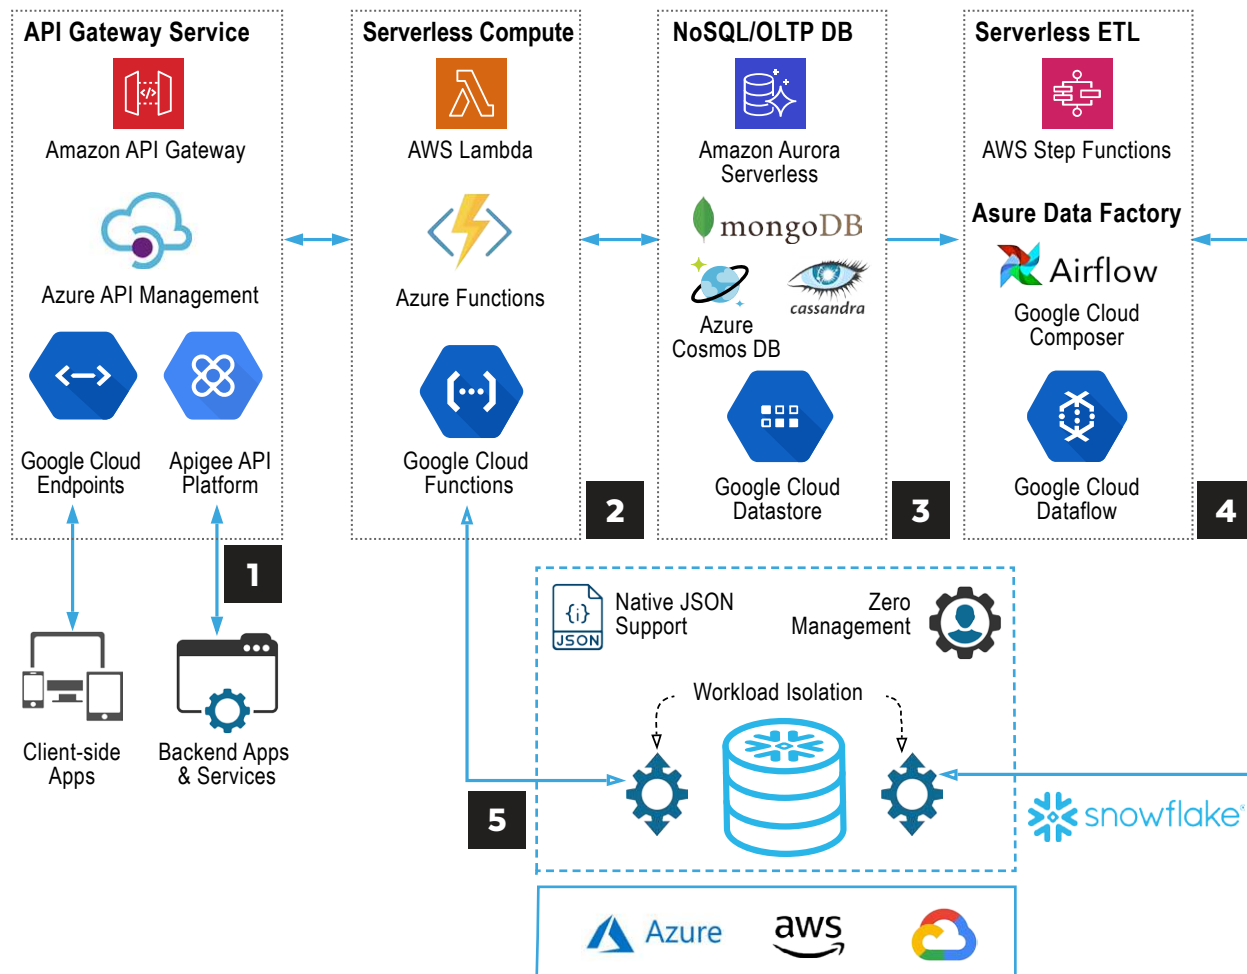


SERVERLESS DATA STACK REFERENCE ARCHITECTURE

SERVERLESS DATA STACK



OBJECTIVE

Build data intensive applications that run on serverless infrastructures.

DESCRIPTION

- 1 The client-side app, running on mobile or web devices, invokes the application logic on the serverless compute via an API gateway service. The gateway authenticates the API calls and throttles them, based on SLAs.
- 2 Serverless compute runs the application logic and scales on demand, without the need to provision or manage servers. The application queries Snowflake data (5) for runtime decisions, such as delivering product recommendations or powering a dashboard for analysis.
- 3 An OLTP or NoSQL database provides the application with high-capacity transaction processing. This NoSQL/OLTP database can also be a serverless service.
- 4 An ETL serverless stack orchestrates the workflow and loads transaction data into Snowflake.
- 5 Snowflake ingests data in batches or in streams and makes it available to the application for queries. Snowflake scales automatically to keep pace with the data pipeline and ensure data is always fresh. Workloads are isolated in virtual warehouses where they can run and scale concurrently without resource contention. Native JSON support enables easy ingestion and querying of flexible schema data alongside structured data.