ClientOptions

CQL

New projects should use CQL. Sending CQL over the wire is not hard, but clients here provide connection pooling, failover, and often other benefits like tracing support.

The list of drivers is in the documentation here. Client authors, please send a pull request to include yours.

Thrift

For older Thrift clients, see ClientOptionsThrift.

Internal API

The StorageProxy API is available to JVM-based clients, but unless you really know that you need it you should probably be using a higher-level client listed above or, failing that, the Thrift API. The StorageProxy API is intended for internal use, and highly specialized use-cases.

Spark

A third party spark connector is available.

Hadoop

Running Hadoop map/reduce jobs in Cassandra is described in HadoopSupport.

https://c.statcounter.com/9397521/0/fe557aad/1/|stats