Can you give an overview of Qube's architecture from a workflow standpoint?

Yes, here is a sample workflow that showcases Qube's main components:

1. An artist submits a job from any Client machine (through the QubeGUI, in-application submission, command-line, python, etc)

2. This creates a package of information (strings, numbers, etc), that is sent to the Supervisor and stored in the database. That package is called a "job."

3. The Supervisor identifies available Workers to process the job.

4. The Supervisor sends the job package to the Worker(s).

5. The Worker service then launches the respective backend (script or executable) that reads the job package and launches the appropriate command line or executable for the rendering.

6. The application (e.g. Maya) then reads in the scene (stored in a central location) and then renders the resulting frames to a central location (like a NAS or other file server).

7. The artist or anyone else, can view the current status of a job through the QubeGUI, command-line, python, etc.