Log Files

Qube! produces a number of different log files, each serving a different purpose. These log files may be stored locally, on a shared file system or both. The log files are an important means to troubleshoot problems, and they serve to document executed jobs.

You should make sure you have enough disk capacity to properly store logs. During the course of normal operation, the Supervisor will produce extensive logs of its activities. By default, these logs are stored in a single file on the Supervisor host called **supelog**. On each Worker, a similar log file is produced in **workerlog**.

For each job, a numbered directory is created corresponding to the job ID. Inside this directory a number of different log files are created containing the history of the job and subjobs, the output from all subjobs, and an archive of the job itself. This is referred to as the **joblog**. This directory and associated log files are created jointly by the Supervisor and the Worker and the Supervisor relays the contents of the logs to various client commands, so they need to be shared in some way. Locating the job logs on the network filesystem will give a meaningful performance improvement, and is highly recommended.

Accessing Log Files

Log files can be most easily accessed from the WranglerView UI, under the "Administration" menu item. On the Supervisor or Worker, you can access the logs directly:

- View Supervisor Log
- View Worker Log

Each gives the option of viewing the entire log, or just the last 250 lines. The number 250 is configurable - see WranglerView Preferences

For Job Logs, select the job, right-click, and then select either:

- Open Job Logs Dir
- Zip and Open Job Logs Dir

This will open a file browser in the directory where the logs for that job are stored.

Other Log Topics

- Job logs
- Supervisor Logs
- Worker Logs
- Writing Job Logs to a Network Filesystem
- Log Rotation
- The logrotate.py script