

Writing Job Logs to a Network Filesystem

This is the quickest and most effective way to decrease the load on your supervisor. If you only change one thing on your farm from the defaults, make this change.



Windows users: This is a bit tricky to set up with a Windows supervisor and workers. You need to ensure that the worker and supervisor services can access the shared filesystem, which usually means having the shared log directory near the top of a share, and having the log directory itself set to Everyone [Full Control]



As well, the paths supplied in `*_logpath` must be UNC, as drive letters are not visible to Windows system services.

Default behavior

Writing logs

The default is to store the logs on the supervisor's local disk. Job log information is handled with **remote log transmission** that follows the following process:

1. The job logs are first stored locally on the Worker
2. Then transmitted from the Worker to the Supervisor
3. Then finally written locally on the Supervisor's filesystem.

Reading logs

When logs are stored on the supervisor, remote log transmission handling follows this protocol:

1. The client asks the supervisor for the logs for a particular job
2. A supervisor process:
 - a. reads the log data for that job into memory
 - b. converts it to a serialized object
 - c. sends that object across a network socket to the client

However, the most efficient way is for both the Supervisor and the Worker to share the job log files directly on a common file server mounted by both the Supervisor host and the Worker hosts. In either case, the Supervisor will need to have access to the entire job log directory structure.

Similarly the Client should read the job log files direct from disk as well instead of having the Supervisor transmit the files to it.

On the Supervisor, job logs will be located in `<supervisor_logpath>/job`. On the Worker, job logs will be located in the `<worker_logpath>/job`. **Both these directories should point to the same location on a shared filesystem.**



The permissions on the shared log directory must be world-writable, which on linux and OS X means `drwxrwxrwx` or mode `0777`, and on Window is Everyone [Full Control]

Steps to Set the Job Log Directory



When you configure the job log directories to be on network storage, you will also need to move the log directories for the older jobs off of the supervisor's local disk up to the network storage; if you skip this step, the logs for those older jobs won't be available.

In order to make this transition as quickly as possible, we recommend that you remove as many jobs as feasible from Qube before starting this reconfiguration.

It's also necessary that no jobs are running while you do this switchover, otherwise you will lose the first portion of that job's logs.

For the supervisor:

Set the Supervisor job log directory to control where the supervisor writes the job logs by modifying the `supervisor_logpath` entry in the `supervisor's qb.conf`:

```
supervisor_logpath = <shared directory>
```

then restart the supervisor service for the change to take effect.

For the workers:

Set the Worker job log directory to control where the supervisor writes the job logs by modifying both the `worker_logpath` and `worker_logmode` entries in either the `qbwrc.conf` (recommended) or each worker's `qb.conf`:

```
worker_logmode = mounted
worker_logpath = <shared directory>
```

If you make the changes in the `qbwrc.conf` on the supervisor, push the changes out with `qbadm w --reconfigure`. (See: [Centralized Worker Configuration](#)). If you edited each worker's `qb.conf`, you will need to restart the worker service for the change to take effect.

For the clients:

Set the Client job log directory. Modify the `client_logpath` entry in each client machine's `qb.conf` so the client machines will directly access the job log files from disk instead of going through the Supervisor:

```
client_logpath = <shared directory>
```



If logs are being written to network storage **but the client does not have the `client_logpath` parameter set** to point to the root of the log directories, the client **will not be able to access the logs** directly and **will fallback to retrieving the logs from the supervisor**.

To test:

- Submit a new job that is very simple, perhaps one that only runs the "set" command. You just want a job that starts, prints a few lines, and exits.
- Verify the job log directory is being created in the expected location
 - If not, the supervisor is not set correctly. Verify and correct, restart the supervisor service, and re-submit to test.
- Verify that the job log directory contains at minimum a `.qja` and `.xja` file. These are written by the supervisor.
- Once the job is complete, verify that the job log directory contains at least a `.out` file (there will be 1 per job instance). There should probably also be a `.err` file. These are written by the worker.
 - If no `.out` or `.err` files exist, or the `.out` does not contain anything that looks like it came from the job itself, then the workers' `logmode` and `logpath` are not set correctly. Verify and correct, the re-submit to test.