Job logs

Logs for each job are stored in subdirectories within the logpath directory. Within each job directory, there are a number of log files associated with the job and its subjobs. The jobs are located in groups of 1000, where jobs 1-999 go in group 0, jobs 1000-1999 go in group 1, etc. The logs are grouped this way to make archiving simple. The exact specification is:

`$LOGPATH/job/`*`job_ID_1000/job_ID`*

Example

- Linux   `/var/spool/qube/job/5000/5042`
- Windows Vista/2008/7 `C:\ProgramData\pfx\qube\logs\job\5000\5042`
- Windows XP/2003   `C:\Program Files\pfx\qube\logs\job\5000\5042`

Contents

The naming convention for each job log file is:
   job_ID.type
where the *job_ID* is the ID of the job, and *type* is the type of log file. The type of a job log file can be either arc for a job archive, or hst for a job history file.

Similarly, the naming convention for each subjob file is:
   *job_ID{}*subjob_ID.type_
where *job_ID* is the ID of the job, *subjob_ID* is the ID of the subjob, and *type* is the type of log file. The type of the log file can be out for standard output, err for standard error output, or sts for job status information. Below is a table summarizing the various naming conventions for log files.

Standard Output	*job_ID{}*subjob_ID{_}.out
Standard Error	*job_ID_ subjob_ID*.err
Job History Log	*job_ID*.hist
Job Statistics	*job_ID_ subjob_ID*.sts
Job Archive File	*job_ID*.qja
Job XML Archive File	*job_ID*.xja
Job Account File	*job_ID*.acc
Job Callback Log	*job_ID*.cb

All but the binary job archive files are human-readable. Additionally, the output and error logs can be can be accessed from the command line tools qbout and qberr.

Verify that the Supervisor and Workers can access the appropriate log file directory, and that the directory permissions are set correctly.

> ⚠️ **Note**
> Under normal circumstances, the Supervisor will automatically create a job log subdirectory when it registers a submitted job. If the Supervisor is unable to create the directory, job execution may fail. Also, since the Workers are responsible for writing output to the subjob log files, if those files cannot be created or written to, job execution may also fail.