

Conversion from McFarm 8.5 to 9.0

April 2003

The UTA Team

McFarm version 9.0 is now available and ready for release p13.08.00. Here is a brief summary of the changes:

- Integration with PBS and Condor
There is now a way to hand-off a Monte Carlo job to a pbs or Condor batch-queue. McFarm still decides which node to run it on, and the batch-queue (under your control) decides when.
- A “SAM Acquire” daemon
There is now a facility to allow you to re-process files that were previously stored in SAM. Effective with new D0 phase “mcp14”, d0gstar MC files are to be stored in SAM against the possibility of re-processing. McFarm will let you create jobs that use such files for input, retrieve them automatically and spread them out in an intelligent manner to cache directories, run your jobs, and then cleanup the cache.
- Support for “merged” thumbnail files
Effective with p13.08.00, farmers are requested to merge the many small ‘tmb’ files for a request into larger file(s) for actual storage. McFarm supports this process with a daemon to stage the tmb files and another daemon to merge and store them when the request is done.
- Job-priorities
It is now allowed to specify a priority from 0 to 10 (0 = high) when using reg_job. This will impact McFarm’s distribution order, the Linux “nice” number, and any batch-queue processing order.
- Improved gathering and distribution
Better handling of SAM error conditions, smoother distribution algorithm.
- Automatic error-out of looping jobs.

If you have McFarm 8.5 configured, this document is a shortcut to the new features and configuration changes you must do. This is followed by a summary of how to integrate with certain batch queues.

- Do a **stop_farm**
This is required anytime new software is to be loaded.
- Obtain recent software
 - Get McFarm 9.0 from our website (www-hep.uta.edu/~d0race/McFarm/McFarm.html).
 - Make a backup just in case:

```
cd /home/mcfarm
tar cf save85.tar bin/* python/* conf_files/*
```
 - Perform an untar into your mcfarm directory:

```
tar xzf mcfarm_08_05.tgz.
```
- 4. Obtain new mc_runjob and new D0 minitars/releases if you wish (not required).
- Modify the bin/setup_farm configuration file. (See the setup_farm.template for examples).
 - You may remove the line containing FARM_JOB_PRIORITY, since priority is now a job-by-job specification.
 - When you are ready to use the new “acquire” daemon, there are two new variables to control its operation: FARM_SAM_MAX_ACQUIRE_FILES and FARM_SAM_ACQUISITION_GROUP.
 - In order to support merging of thumbnail files, you must create a subdirectory in any file-server cache called “stage”, and then add the following merge parameters:

```
export FARM_MERGE_STAGE_DIR=$FARM_FILESERVER_CACHE_001_A/stage
export FARM_MERGE_STAGE_NODE=001
export FARM_MERGE_INTERVAL_SECS=300
```

- You may want to add lines for FARM_LOOP_MINS and FARM_LOOP_ACTION to tell McFarm how long to wait before declaring a job to be looping, and what to do about it.

IMPORTANT: After saving the setup_farm script, don't forget to re-issue “. bin/setup_farm” in all active login sessions!

- Modify the distribute.conf configuration file:

- You may remove lines at the end that are associated with job-selection, such as “gather_list=”, “job=”, and “allownode=”. The ability to control the order of distribution is now handled by job priorities (along with the order of jobs in the gather.conf file). The distribute daemon will now examine the gather.conf file (and the contents of the dist_queue) for job selections.
- If you will be using Condor or PBS batch queues, see the user's guide for modifications to node-configuration that are done to this file.

- Modify your “start_farm” script

- You will want to add a line for “start_gather -m=MERGE” in order to auto-start the two merge daemons (part A – the staging, and part B – the merging)

- Create a new link on each node

On each node, McFarm now requires a numbered-link to the local scratch disk. As root, perform these commands on each node (using your directory-prefix if it is not already ‘scr’):

```
ln -s /scratch /scrNNN (where NNN is that node number)
chown mcfarm /scrNNN
```

- Do a **start_farm** after doing a “. bin/setup_farm”

Ensure that all daemons come back up without messages before proceeding.

- Alter your future job-creation parameters

- Instead of just sending d0gstar files to “metadata”, farmers are now requested to actually store them in SAM, so use “—dogstar_disposition=sam” in all d0gstar jobs.
- Thumbnail files are not to be stored directly, instead use “—thumbnail_disposition=merge”, which will declare the tmb files but stage them (using the merge_a daeon) for later storage (by the merge_b daemon).