

Table of Contents

Job Recovery.....	1
Introduction.....	2
Algorithm.....	3
Job Recovery Error Codes.....	5

Job Recovery

Introduction

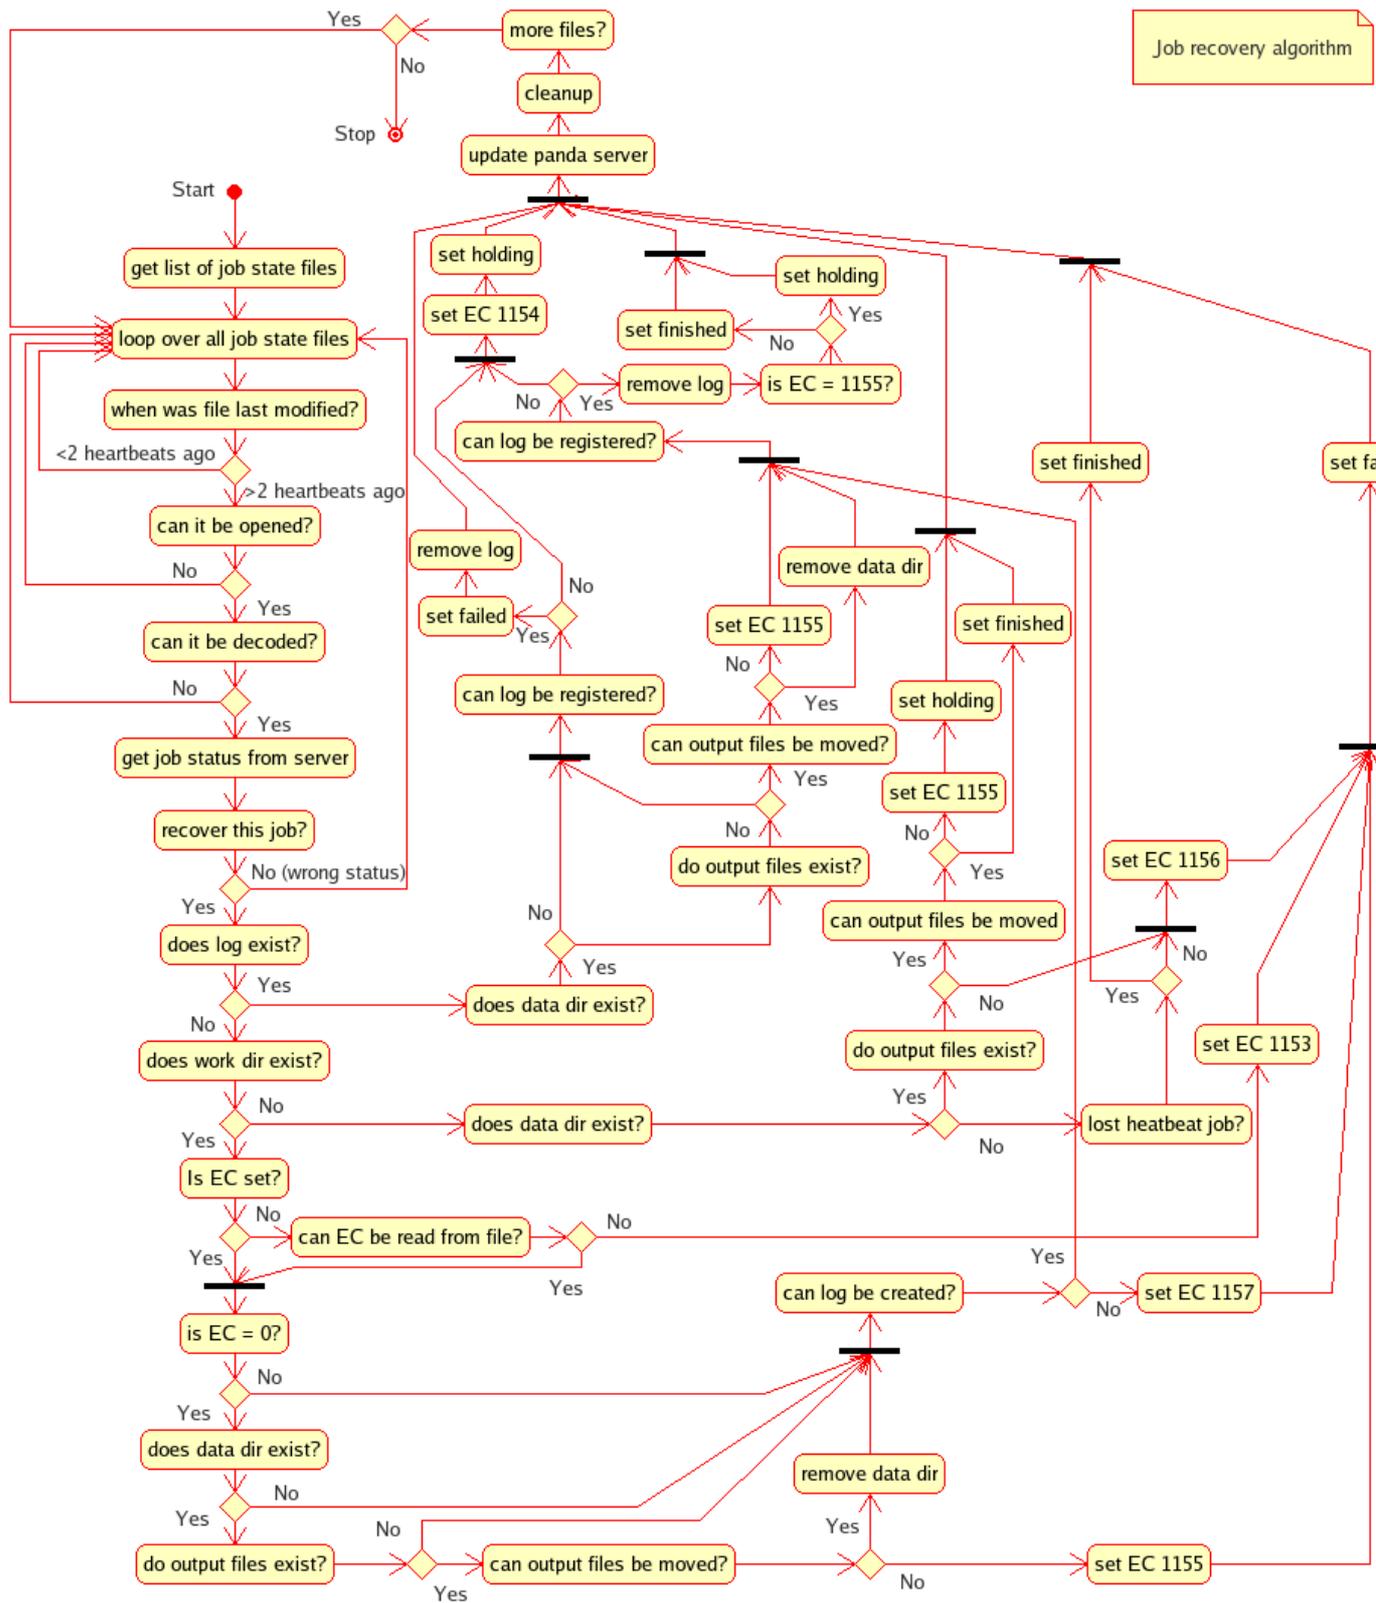
The task of the Panda pilot job recovery is to scan the disk for old 'lost' job directories and log files. When it finds a lost job, in the form of a job state file which has not been modified recently, it will try to recover the job and update the Panda server for it.

Algorithm

The Panda pilot job recovery works by scanning all work directories found on the local disk for existing job state files. When such a file is found, the algorithm first checks when the job state file was last updated; if it was less than two heartbeats ago (i.e. 60 minutes), it will skip it temporarily (it will be picked up again at a later time by the next pilot arriving to the node). If the last update was done more than two heartbeats ago, it will try to read the file. If the work directory of the lost job still exists, it will try to create the log file. If it doesn't exist, it will look for the compressed work directory, i.e. the log. If it finds anything, it will register the log. If there are no work directories or log files, the job state will be set to failed with an pilot exit code of 1153 (lost job did not finished). Any existing remaining data files (from a failed DQ2 transfer) will be moved to the local SE and registered. Before ending the iteration of the job state file loop, the Panda server will be updated with the now known fate of the lost job. The work directory will be cleaned for successful recoveries (also if the lost job can never be recovered). Full implementation details can be see in the activity diagram below.

The algorithm is executed at the pilot startup.

The lost job recovery can be switched on/off with the "-j True/False" pilot flag. Job state files will still be maintained.



Activity diagram for the job recovery algorithm. [Draft 12, implemented in pilot v SPOCK 1f]

Job Recovery Error Codes

Error Code	Description
1153	Job Recovery: Lost job did not finish - lost forever
1154	Job Recovery: Failed to re-register log file
1155	Job Recovery: Failed to move output files
1156	Job Recovery: Missing log and work dir
1157	Job Recovery: Could not create log

Major updates:

-- PaulNilsson - 06 Oct 2006

Responsible: PaulNilsson

This topic: PanDA > PandaPilotJobRecovery

Topic revision: r12 - 2006-12-21 - PaulNilssonSecondary



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback