

Hjobkill

#FORMAT plain hjobkill(1) HADOOP hjobkill(1)

NAME

hjobkill - Kill jobs given the job name and clustername. it maybe rerun depending on parameter.

SYNOPSIS

hjobkill [OPTION] jobid1 jobid2 ...

DESCRIPTION

hjobkill will kill the jobs specified by the parameter. If 0 is the argument, all jobs will be terminated.

-u user_name

With this option, hjobs will display jobs for a user_name

-jt jobtracker url link.

specify the cluster cluster by the jobtracker

url link, for example. -jt hostA:3479;hostB:3479

-r

If this option is present, the job will be reschedule and ran.

EXAMPLE

% hjobkill job_200708312310_0001

Job <job_200709181000_0001 @c0101:8802> has been killed successfully.

% hjobkill -r job_200708312310_0001

Job <job_200709181000_0001 @c0101:8802> has been killed successfully.

Job <job_200709181000_0001 @c0101:8802> has been rerun.

% hjobkill 0

Job <job_200709181000_0001 @c0101:8802> has been killed successfully.

Job <job_200709181000_0002 @c0101:8802> has been killed successfully.

Job <job_200709181000_0003 @c0101:8802> has been killed successfully.

AUTHOR

Written by Shuguang Liu

REPORT BUGS

Report bugs to <honkliu@hotmail.com>

COPYRIGHT

Copyright (C) 2005-2010 Shuguang Liu's Inc.

This file is not a free software, all the source codes are protected and will not be released to any body or organization without authority.

SEE ALSO

hjobs(1)

hjobkill (hadoop utils) 13.0 October 2007