



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LCRM psub to Moab msub Translation Guide

Donald A. Lipari

January 25, 2010

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory in part under Contract W-7405-Eng-48 and in part under Contract DE-AC52-07NA27344.

The psub wrapper support is subject to the following scenarios (annotated in column 4 in the table below).

1. The option is accepted and submitted as a specification of a Moab job (accept).
2. The option is ignored and a warning is issued (ignore & warn).
3. The option is archaic or irrelevant and ignored. A warning is seen only when psub -v is present (ignore).
4. The option is not supported and the submission is rejected (reject).

psub Option	Argument	Description	psub Support	msub Option	srun Argument
-H		help	accept	--help	N/A
-A	date-time	run at time	accept	-a	N/A
-b	bank	bank	accept	-A	N/A
-bgl	[geometry=X[xY[xZ]]], [rotate norotate], [conn_type=mesh torus nav]	BG/L attributes only	accept on BG/L systems reject otherwise	--slurm ...	--geometry=XxYxZ (if specified) --no-rotate (if -bgl norotate) --conn-type=MESH TORUS NAV
-c	constraint(s)		accept	-l partition=<machine> -l feature=<val>	N/A
-cpn	CPUs per node	(archaic) same as -np	ignore	see -np below	see -np
-d	job-ID	dependent job	accept	-l depend=<job-ID> or -l depend= afterok:<job-ID>	--dependency=<job-ID>
-e	file	error output file	accept	-e	--error
-eo		merge error output into file output	accept	-j oe	N/A
-exempt	exemption list	remove job limits for certain users	ignore (for now)	not yet supported	N/A
-expedite		confers top priority	accept	-l qos=expedite	N/A
-g	geometry	[tasks][switch][@layout]	accept only on AIX	-l nodes=<nodes>: ppn=<taskspernode>	add switch (ip or us) to --network=type string
-i	command	invoke "command" instead of reading from	accept	echo "command" msub	N/A

psub Option	Argument	Description	psub Support	msub Option	srun Argument
		script file			
-lc	limit	core file size limit	ignore & warn		N/A
-ld	limit	data segment limit	ignore & warn		N/A
-lF	limit	file size limit for job	ignore & warn		N/A
-lf	limit	file size limit per process	ignore & warn		N/A
-lM	limit	RSS size expectation	accept	-l dmem=<val>	N/A
-ln	node spec	a complex node specification; see supplemental info to be provided	accept	-l nodes=<val>	--nodes
-lo	limit	open file count limit	ignore & warn		N/A
-lr	limit	per-process maximum resident set size limit	ignore & warn		N/A
-ls	limit	per-process maximum stack size limit	ignore & warn		N/A
-lt	limit	CPU time limit per process	ignore & warn		N/A
-mb		send mail at job start	accept	-m b	--mail-type=BEGIN
-me		send mail at job completion	accept	-m e	--mail-type=END
-mn		suppress mail from RM	ignore	-m n	N/A
-nettype	type	SN_ALL or SN_SINGLE	ignore	convey via srun arguments	add SN_ALL or SN_SINGLE to --network=type string
-nobulkxfer		turns off bulk transfer	ignore	convey via srun arguments	add BULK_XFER to --network=type string
-nokill		keep job running if a node fails	accept	-l resfailpolicy=ignore	--no-kill
-np	number of procesors N/A for parallel clusters	requested processors on single node SMP	accept	-l ttc=<val>	--ntasks

psub Option	Argument	Description	psub Support	msub Option	srun Argument
	(job gets whole node).	machines			
-nr		do not re-queue job after node failure	ignore	-l resfailpolicy=cancel	--no-requeue
N/A		re-queue job after node failure	ignore	-l resfailpolicy=requeue	--requeue
-o	file	output file	accept	-o	--output
-p	priority	gives user the power to elevate their job's priority	accept	-p	N/A
-pool	SLURM partition	targets job to a specific node partition	accept	-q [<queue>] [@<server>]	--partition
-prj	project name	user specified comment field – not tracked	ignore	-l var:Project=<name>	N/A
-r	job name	user specified name for their job	accept	-N	--job-name
-s	shell name	absolute path name	accept	-S	N/A
-standby		can be pre-empted	accept	-l qos=standby	N/A
-tM	time	CPU time limit	accept	use -tW	N/A
-tW	time	wall-clock time limit	accept	-l walltime=<val>	--time
-v		provide verbose psub output	accept	N/A	N/A
-x		use current environment variables at job execution	accept	-V	N/A
Obsolete Arguments - Ignore					
-creds	credentials	Kerberos support	ignore		N/A
-dm		never implemented	ignore		N/A
-ke		leave output file on execution host	ignore		N/A
-ko		leave output file on execution host	ignore		N/A
-nc		non-checkpointable	ignore		N/A
-net	protocol	mpi or lapi	ignore		N/A
-re		flush error output	ignore		N/A
-ro		flush output file	ignore		N/A

psub accepts a job command file argument of the user's script. The first portion of that script contains #PSUB lines that contain psub arguments. The remaining portion of the script contains the body of the user's job.

msub also accepts a job command file argument of the user's Moab-compliant script. The msub options must also be listed in the first portion of the script and must be prefixed with "#MSUB". Note that, unlike LCRM (which allows a space between the # and PSUB), Moab does not accept a space between the # and the MSUB.

A utility is available in /usr/bin (/opt/freeware/bin on AIX), *lcrm2moab*, that converts a legacy LCRM job command script to the msub-style format.