

*(Click to open topic with navigation)*

## qsub

Submit PBS job.

---

### Synopsis

```
qsub [-a date_time] [-A account_string] [-b secs] [-c checkpoint_options]
[-C directive_prefix] [-d path] [-D path] [-e path] [-f] [-F] [-h]
[-I ] [-j join ] [-k keep ] [-l resource_list ]
[-m mail_options] [-M user_list] [-n] [-N name] [-o path]
[-p priority] [-P user[:group]] [-q destination] [-r c] [-S path_to_shell(s)]
[-t array_request] [-u user_list]
[-v variable_list] [-V] [-W additional_attributes] [-x] [-X] [-z] [script]
```

---

### Description

To create a job is to submit an executable script to a batch server. The batch server will be the default server unless the **-q** option is specified. The command parses a script prior to the actual script execution; it does not execute a script itself. All script-writing rules remain in effect, including the **"#!"** at the head of the file (see discussion of PBS\_DEFAULT under [Environment variables](#)). Typically, the script is a shell script which will be executed by a command shell such as sh or csh.

Options on the qsub command allow the specification of attributes which affect the behavior of the job.

The qsub command will pass certain environment variables in the Variable\_List attribute of the job. These variables will be available to the job. The value for the following variables will be taken from the environment of the qsub command: HOME, LANG, LOGNAME, PATH, MAIL, SHELL, and TZ. These values will be assigned to a new name which is the current name prefixed with the string "PBS\_O\_". For example, the job will have access to an environment variable named PBS\_O\_HOME which have the value of the variable HOME in the qsub command environment.

In addition to the above, the following environment variables will be available to the batch job:

Variable	Description
<b>PBS_O_HOST</b>	The name of the host upon which the qsub command is running.
<b>PBS_SERVER</b>	The hostname of the pbs_server which qsub submits the job to.
<b>PBS_O_QUEUE</b>	The name of the original queue to which the job was submitted.

Variable	Description
<b>PBS_O_WORKDIR</b>	The absolute path of the current working directory of the qsub command.
<b>PBS_ARRAYID</b>	Each member of a job array is assigned a unique identifier (see <b>-t</b> option).
<b>PBS_ENVIRONMENT</b>	Set to PBS_BATCH to indicate the job is a batch job, or to PBS_INTERACTIVE to indicate the job is a PBS interactive job (see <b>-I</b> option).
<b>PBS_JOBID</b>	The job identifier assigned to the job by the batch system. It can be used in the stdout and stderr paths. TORQUE replaces \$PBS_JOBID with the job's jobid (for example, #PBS -o /tmp/\$PBS_JOBID.output).
<b>PBS_JOBNAME</b>	The job name supplied by the user.
<b>PBS_NODEFILE</b>	The name of the file contain the list of nodes assigned to the job (for parallel and cluster systems).
<b>PBS_QUEUE</b>	The name of the queue from which the job is executed.

---

## Options

Option	Name	Description
<b>-a</b>	date_time	<p>Declares the time after which the job is eligible for execution. The date_time argument is in the form:</p> <pre>[[[CC]YY]MM]DD]hhmm[.SS]</pre> <p>where <i>CC</i> is the first two digits of the year (the century), <i>YY</i> is the second two digits of the year, <i>MM</i> is the two digits for the month, <i>DD</i> is the day of the month, <i>hh</i> is the hour, <i>mm</i> is the minute, and the optional <i>SS</i> is the seconds.</p> <p>If the month (MM) is not specified, it will default to the current month if the specified day (DD) is in the future. Otherwise, the month will be set to next month. Likewise, if the day (DD) is not specified, it will default to today if the time (hhmm) is in the future. Otherwise, the day will be set to tomorrow.</p> <p>For example, if you submit a job at 11:15 am with a time of <b>-a 1110</b>, the job will be eligible to run at 11:10 am tomorrow.</p>
<b>-A</b>	account_string	<p>Defines the account string associated with the job. The account_string is an undefined string of characters and is interpreted by the server which executes the job. See section 2.7.1 of the PBS ERS.</p>
<b>-b</b>	seconds	<p>Defines the maximum number of seconds qsub will block attempting to contact pbs_server. If pbs_server is down, or for a variety of communication failures, qsub will continually retry</p>

Option	Name	Description
		<p>connecting to pbs_server for job submission.</p> <p>This value overrides the CLIENTRETRY parameter in <code>torque.cfg</code>. This is a non-portable TORQUE extension. Portability-minded users can use the PBS_CLIENRETRY environmental variable. A negative value is interpreted as infinity. The default is 0.</p>
<b>-c</b>	checkpoint_options	<p>Defines the options that will apply to the job. If the job executes upon a host which does not support checkpoint, these options will be ignored.</p> <p>Valid checkpoint options are:</p> <ul style="list-style-type: none"> <li>• <b>none</b> – No checkpointing is to be performed.</li> <li>• <b>enabled</b> – Specify that checkpointing is allowed but must be explicitly invoked by either the <b>qhold</b> or <b>qchkpt</b> commands.</li> <li>• <b>shutdown</b> – Specify that checkpointing is to be done on a job at pbs_mom shutdown.</li> <li>• <b>periodic</b> – Specify that periodic checkpointing is enabled. The default interval is 10 minutes and can be changed by the <code>\$checkpoint_interval</code> option in the MOM config file or by specifying an interval when the job is submitted</li> <li>• <b>interval=minutes</b> – Checkpointing is to be performed at an interval of minutes, which is the integer number of minutes of wall time used by the job. This value must be greater than zero.</li> <li>• <b>depth=number</b> – Specify a number (depth) of checkpoint images to be kept in the checkpoint directory.</li> <li>• <b>dir=path</b> – Specify a checkpoint directory (default is <code>/var/spool/torque/checkpoint</code>).</li> </ul>
<b>-C</b>	directive_prefix	<p>Defines the prefix that declares a directive to the qsub command within the script file. (See the paragraph on script directives under <a href="#">Extended description</a>.)</p> <p>If the <code>-c</code> option is presented with a <code>directive_prefix</code> argument that is the null string, qsub will not scan the script file for directives.</p>
<b>-d</b>	path	<p>Defines the working directory path to be used for the job. If the <code>-d</code> option is not specified, the default working directory is the home directory. This option sets the environment variable <code>PBS_O_INITDIR</code>.</p>
<b>-D</b>	path	<p>Defines the root directory to be used for the job. This option sets the environment variable <code>PBS_O_ROOTDIR</code>.</p>
<b>-e</b>	path	<p>Defines the path to be used for the standard error stream of the batch job. The path argument is of the form:</p> <p><code>[hostname:]path_name</code></p> <p>where <i>hostname</i> is the name of a host to which the file will be returned, and <i>path_name</i> is the path name on that host in the syntax recognized by POSIX. The argument will be interpreted as</p>

Option	Name	Description
		<p>follows:</p> <ul style="list-style-type: none"> <li>• <b>path_name</b> – where <i>path_name</i> is not an absolute path name, then the qsub command will expand the path name relative to the current working directory of the command. The command will supply the name of the host upon which it is executing for the hostname component.</li> <li>• <b>hostname:path_name</b> – where <i>path_name</i> is not an absolute path name, then the qsub command will not expand the path name relative to the current working directory of the command. On delivery of the standard error, the path name will be expanded relative to the users home directory on the hostname system.</li> <li>• <b>path_name</b> – where <i>path_name</i> specifies an absolute path name, then the qsub will supply the name of the host on which it is executing for the hostname.</li> <li>• <b>hostname:path_name</b> – where <i>path_name</i> specifies an absolute path name, the path will be used as specified.</li> </ul> <p>If the <code>-e</code> option is not specified, the default file name for the standard error stream will be used. The default name has the following form:</p> <ul style="list-style-type: none"> <li>• <b>job_name.esquence_number</b> – where <i>job_name</i> is the name of the job (see the <code>-n</code> name option) and <i>sequence_number</i> is the job number assigned when the job is submitted.</li> </ul>
<b>-f</b>	---	<p>Job is made fault tolerant. Jobs running on multiple nodes are periodically polled by mother superior. If one of the nodes fails to report, the job is canceled by mother superior and a failure is reported. If a job is fault tolerant, it will not be canceled based on failed polling (no matter how many nodes fail to report). This may be desirable if transient network failures are causing large jobs not to complete, where ignoring one failed polling attempt can be corrected at the next polling attempt.</p> <div style="border: 1px solid black; border-radius: 10px; padding: 5px; margin-top: 10px;"> <p><b>i</b> If TORQUE is compiled with <code>PBS_NO_POSIX_VIOLATION</code> (there is no config option for this), you have to use <code>-w fault_tolerant=true</code> to mark the job as fault tolerant.</p> </div>
<b>-F</b>	---	<p>Specifies the arguments that will be passed to the job script when the script is launched. The accepted syntax is:</p> <pre>qsub -F "myarg1 myarg2 myarg3=myarg3value" myscript2.sh</pre> <div style="border: 1px solid black; border-radius: 10px; padding: 5px; margin-top: 10px;"> <p><b>i</b> Quotation marks are required. qsub will fail with an error message if the argument following <code>-F</code> is not a quoted value. The pbs_mom server will pass the quoted value as arguments to the job script when it launches the script.</p> </div>
<b>-h</b>	---	<p>Specifies that a user hold be applied to the job at submission time.</p>

Option	Name	Description
<b>-I</b>	---	Declares that the job is to be run "interactively". The job will be queued and scheduled as any PBS batch job, but when executed, the standard input, output, and error streams of the job are connected through qsub to the terminal session in which qsub is running. Interactive jobs are forced to not rerunable. See <a href="#">Extended description</a> for additional information of interactive jobs.
<b>-j</b>	join	<p>Declares if the standard error stream of the job will be merged with the standard output stream of the job.</p> <p>An option argument value of <b>oe</b> directs that the two streams will be merged, intermixed, as standard output. An option argument value of <b>eo</b> directs that the two streams will be merged, intermixed, as standard error.</p> <p>If the join argument is <b>n</b> or the option is not specified, the two streams will be two separate files.</p>
<b>-k</b>	keep	<p>Defines which (if either) of standard output or standard error will be retained on the execution host. If set for a stream, this option overrides the path name for that stream. If not set, neither stream is retained on the execution host.</p> <p>The argument is either the single letter "e" or "o", or the letters "e" and "o" combined in either order. Or the argument is the letter "n".</p> <ul style="list-style-type: none"> <li>• <b>e</b> – The standard error stream is to retained on the execution host. The stream will be placed in the home directory of the user under whose user id the job executed. The file name will be the default file name given by:  <code>job_name.e<sub>sequence</sub></code>  where <i>job_name</i> is the name specified for the job, and <i>sequence</i> is the sequence number component of the job identifier.</li> <li>• <b>o</b> – The standard output stream is to retained on the execution host. The stream will be placed in the home directory of the user under whose user id the job executed. The file name will be the default file name given by:  <code>job_name.o<sub>sequence</sub></code>  where <i>job_name</i> is the name specified for the job, and <i>sequence</i> is the sequence number component of the job identifier.</li> <li>• <b>eo</b> – Both the standard output and standard error streams will be retained.</li> <li>• <b>oe</b> – Both the standard output and standard error streams will be retained.</li> <li>• <b>n</b> – Neither stream is retained.</li> </ul>
<b>-l</b>	resource_list	<p>Defines the resources that are required by the job and establishes a limit to the amount of resource that can be consumed. If not set for a generally available resource, such as CPU time, the limit is infinite. The resource_list argument is of the form:</p> <pre>resource_name[=[value]][,resource_name[=[value]],...]</pre>

Option	Name	Description
		<div style="border: 1px solid black; border-radius: 10px; padding: 10px; margin-bottom: 10px;"> <p><b>i</b> In this situation, you should request the more inclusive resource first. For example, a request for procs should come before a gres request.</p> </div> <p>For information on specifying multiple types of resources for allocation, see "Multi-Req Support" under "General Job Policies" in the <a href="#">Moab Workload Manager</a> documentation.</p>
<b>-m</b>	mail_options	<p>Defines the set of conditions under which the execution server will send a mail message about the job. The mail_options argument is a string which consists of either the single character "n", or one or more of the characters "a", "b", and "e".</p> <p>If the character "n" is specified, no normal mail is sent. Mail for job cancels and other events outside of normal job processing are still sent.</p> <p>For the letters "a", "b", and "e":</p> <ul style="list-style-type: none"> <li>• <b>a</b> – Mail is sent when the job is aborted by the batch system.</li> <li>• <b>b</b> – Mail is sent when the job begins execution.</li> <li>• <b>e</b> – Mail is sent when the job terminates.</li> </ul> <p>If the <code>-m</code> option is not specified, mail will be sent if the job is aborted.</p>
<b>-M</b>	user_list	<p>Declares the list of users to whom mail is sent by the execution server when it sends mail about the job.</p> <p>The user_list argument is of the form:</p> <pre>user[@host][,user[@host],...]</pre> <p>If unset, the list defaults to the submitting user at the qsub host, i.e. the job owner.</p>
<b>-n</b>	node-exclusive	<p>Allows a user to specify an exclusive-node access/allocation request for the job. This affects only cpusets and compatible schedulers (see <a href="#">Linux cpuset support</a>).</p>
<b>-N</b>	name	<p>Declares a name for the job. The name specified may be an unlimited number of characters in length. It must consist of printable, non white space characters with the first character alphabetic.</p> <p>If the <code>-N</code> option is not specified, the job name will be the base name of the job script file specified on the command line. If no script file name was specified and the script was read from the standard input, then the job name will be set to STDIN.</p>
<b>-o</b>	path	<p>Defines the path to be used for the standard output stream of the batch job. The path argument is of the form:</p> <pre>[hostname:]path_name</pre> <p>where <i>hostname</i> is the name of a host to which the file will be</p>

Option	Name	Description
		<p>returned, and <i>path_name</i> is the path name on that host in the syntax recognized by POSIX. The argument will be interpreted as follows:</p> <ul style="list-style-type: none"> <li>• <b>path_name</b> – where <i>path_name</i> is not an absolute path name, then the qsub command will expand the path name relative to the current working directory of the command. The command will supply the name of the host upon which it is executing for the hostname component.</li> <li>• <b>hostname:path_name</b> – where <i>path_name</i> is not an absolute path name, then the qsub command will not expand the path name relative to the current working directory of the command. On delivery of the standard output, the path name will be expanded relative to the users home directory on the hostname system.</li> <li>• <b>path_name</b> – where <i>path_name</i> specifies an absolute path name, then the qsub will supply the name of the host on which it is executing for the hostname.</li> <li>• <b>hostname:path_name</b> where <i>path_name</i> specifies an absolute path name, the path will be used as specified.</li> </ul> <p>If the <code>-o</code> option is not specified, the default file name for the standard output stream will be used. The default name has the following form:</p> <ul style="list-style-type: none"> <li>• <b>job_name.osequence_number</b> – where <i>job_name</i> is the name of the job (see the <code>-n</code> name option) and <i>sequence_number</i> is the job number assigned when the job is submitted.</li> </ul>
<b>-p</b>	priority	<p>Defines the priority of the job. The priority argument must be a integer between -1024 and +1023 inclusive. The default is no priority which is equivalent to a priority of zero.</p>
<b>-P</b>	user[:group]	<p>Allows a root user to submit a job as another user. TORQUE treats proxy jobs as though the jobs were submitted by the supplied username. This feature is available in TORQUE 2.4.7 and later, however, TORQUE 2.4.7 does not have the ability to supply the <code>[:group]</code> option; it is available in TORQUE 2.4.8 and later.</p>
<b>-q</b>	destination	<p>Defines the destination of the job. The destination names a queue, a server, or a queue at a server.</p> <p>The qsub command will submit the script to the server defined by the destination argument. If the destination is a routing queue, the job may be routed by the server to a new destination.</p> <p>If the <code>-q</code> option is not specified, the qsub command will submit the script to the default server. (See <a href="#">Environment variables</a> and the PBS ERS section 2.7.4, "Default Server".)</p> <p>If the <code>-q</code> option is specified, it is in one of the following three forms:</p> <ul style="list-style-type: none"> <li>• queue</li> <li>• @server</li> <li>• queue@server</li> </ul>

Option	Name	Description
		<p>If the destination argument names a queue and does not name a server, the job will be submitted to the named queue at the default server.</p> <p>If the destination argument names a server and does not name a queue, the job will be submitted to the default queue at the named server.</p> <p>If the destination argument names both a queue and a server, the job will be submitted to the named queue at the named server.</p>
<b>-r</b>	y/n	<p>Declares whether the job is rerunable (see the <b>qrerun</b> command). The option argument is a single character, either y or n.</p> <p>If the argument is "y", the job is rerunable. If the argument is "n", the job is not rerunable. The default value is y, rerunable.</p>
<b>-S</b>	path_list	<p>Declares the path to the desired shell for this job.</p> <pre>qsub script.sh -S /bin/tcsh</pre> <p>The option argument path_list is in the form:</p> <pre>path[@host][,path[@host],...]</pre> <p>Only one path may be specified for any host named. Only one path may be specified without the corresponding host name. The path selected will be the one with the host name that matched the name of the execution host. If no matching host is found, then the path specified without a host will be selected, if present.</p> <p>If the <b>-s</b> option is not specified, the option argument is the null string, or no entry from the path_list is selected, the execution will use the user's login shell on the execution host.</p>
<b>-t</b>	array_request	<p>Specifies the task ids of a job array. Single task arrays are allowed. The array_request argument is an integer id or a range of integers. Multiple ids or id ranges can be combined in a comma delimited list. Examples: <b>-t 1-100</b> or <b>-t 1,10,50-100</b></p> <p>An optional <b>slot limit</b> can be specified to limit the amount of jobs that can run concurrently in the job array. The default value is unlimited. The slot limit must be the last thing specified in the array_request and is delimited from the array by a percent sign (<b>%</b>).</p> <pre>qsub script.sh -t 0-299%5</pre> <p>This sets the slot limit to 5. Only 5 jobs from this array can run at the same time.</p> <p>You can use <b>qalter</b> to modify slot limits on an array. The server parameter <b>max_slot_limit</b> can be used to set a global slot limit policy. When using slot limits in TORQUE with Moab or Maui, you should also set the <b>moab_array_compatible</b> server parameter to <b>TRUE</b>.</p>
<b>-u</b>	user_list	<p>Defines the user name under which the job is to run on the execution system.</p>



Option	Name	Description
		<p>The <code>user_list</code> argument is of the form:</p> <pre>user[@host][,user[@host],...]</pre> <p>Only one user name may be given per specified host. Only one of the user specifications may be supplied without the corresponding host specification. That user name will be used for execution on any host not named in the argument list. If unset, the user list defaults to the user who is running <code>qsub</code>.</p>
<b>-v</b>	variable_list	<p>Expands the list of environment variables that are exported to the job.</p> <p>In addition to the variables described in the "Description" section above, <code>variable_list</code> names environment variables from the <code>qsub</code> command environment which are made available to the job when it executes. The <code>variable_list</code> is a comma separated list of strings of the form <code>variable</code> or <code>variable=value</code>. These variables and their values are passed to the job.</p>
<b>-V</b>	---	<p>Declares that all environment variables in the <code>qsub</code> commands environment are to be exported to the batch job.</p>
<b>-W</b>	additional_attributes	<p>The <code>-w</code> option allows for the specification of additional job attributes. The general syntax of <code>-w</code> is in the form:</p> <pre>-W attr_name=attr_value[,attr_name=attr_value...]</pre> <div data-bbox="630 1052 1442 1192" style="border: 1px solid black; border-radius: 10px; padding: 5px; margin: 10px 0;"> <p><b>i</b> If white space occurs anywhere within the option argument string or the equal sign, "=", occurs within an <code>attribute_value</code> string, then the string must be enclosed with either single or double quote marks.</p> </div> <p>PBS currently supports the following attributes within the <code>-w</code> option:</p> <ul style="list-style-type: none"> <li>• <b>depend=dependency_list</b> – Defines the dependency between this and other jobs. The <code>dependency_list</code> is in the form: <pre>type[:argument[:argument...]][,type:argument...]</pre> <p>The argument is either a numeric count or a PBS job id according to type. If argument is a count, it must be greater than 0. If it is a job id and not fully specified in the form <code>seq_number.server.name</code>, it will be expanded according to the default server rules which apply to job IDs on most commands. If argument is null (the preceding colon need not be specified), the dependency of the corresponding type is cleared (unset). For more information, see <a href="#">depend=dependency_list valid dependencies</a>.</p> </li> <li>• <b>group_list=g_list</b> – Defines the group name under which the job is to run on the execution system. The <code>g_list</code> argument is of the form: <pre>group[@host][,group[@host],...]</pre> <p>Only one group name may be given per specified host. Only</p> </li> </ul>

Option	Name	Description
		<p>one of the group specifications may be supplied without the corresponding host specification. That group name will be used for execution on any host not named in the argument list. If not set, the <code>group_list</code> defaults to the primary group of the user under which the job will be run.</p> <ul style="list-style-type: none"> <li>• <b>interactive=true</b> – If the interactive attribute is specified, the job is an interactive job. The <b>-I</b> option is an alternative method of specifying this attribute.</li> <li>• <b>stagein=file_list</b></li> <li>• <b>stageout=file_list</b> – Specifies which files are staged (copied) in before job start or staged out after the job completes execution. On completion of the job, all staged-in and staged-out files are removed from the execution system. The <code>file_list</code> is in the form:  <code>local_file@hostname:remote_file[,...]</code> </li> </ul> <p>regardless of the direction of the copy. The name <code>local_file</code> is the name of the file on the system where the job executed. It may be an absolute path or relative to the home directory of the user. The name <code>remote_file</code> is the destination name on the host specified by <code>hostname</code>. The name may be absolute or relative to the user's home directory on the destination host. The use of wildcards in the file name is not recommended. The file names map to a remote copy program (<code>rcp</code>) call on the execution system in the following manner:</p> <ul style="list-style-type: none"> <li>○ For stagein: <code>rcp hostname:remote_file local_file</code></li> <li>○ For stageout: <code>rcp local_file hostname:remote_file</code></li> </ul> <p>Data staging examples:</p> <pre>-W stagein=/tmp/input.txt@headnode:/home/user /input.txt</pre> <pre>-W stageout=/tmp/output.txt@headnode:/home/user /output.txt</pre> <p>If TORQUE has been compiled with <code>wordexp</code> support, then variables can be used in the specified paths. Currently only <code>\$PBS_JOBID</code>, <code>\$HOME</code>, and <code>\$TMPDIR</code> are supported for stagein.</p> <ul style="list-style-type: none"> <li>• <b>umask=XXX</b> – Sets <code>umask</code> used to create <code>stdout</code> and <code>stderr</code> spool files in <code>pbs_mom</code> spool directory. Values starting with 0 are treated as octal values, otherwise the value is treated as a decimal <code>umask</code> value.</li> </ul>
<b>-X</b>	---	Enables X11 forwarding. The <code>DISPLAY</code> environment variable must be set.
<b>-z</b>	---	Directs that the <code>qsub</code> command is not to write the job identifier assigned to the job to the command's standard output.

**depend=dependency\_list valid dependencies**

Dependency	Description
<code>synccount:count</code>	This job is the first in a set of jobs to be executed at the same time. Count is the number of additional jobs in the set.
<code>syncwith:jobid</code>	This job is an additional member of a set of jobs to be executed at the same time. In the above and following dependency types, jobid is the job identifier of the first job in the set.
<code>after:jobid[:jobid...]</code>	This job may be scheduled for execution at any point after jobs jobid have started execution.
<code>afterok:jobid[:jobid...]</code>	This job may be scheduled for execution only after jobs jobid have terminated with no errors. See the csh warning under <a href="#">Extended description</a> .
<code>afternotok:jobid[:jobid...]</code>	This job may be scheduled for execution only after jobs jobid have terminated with errors. See the csh warning under <a href="#">Extended description</a> .
<code>afterany:jobid[:jobid...]</code>	This job may be scheduled for execution after jobs jobid have terminated, with or without errors.
<code>on:count</code>	This job may be scheduled for execution after count dependencies on other jobs have been satisfied. This form is used in conjunction with one of the "before" forms (see below).
<code>before:jobid[:jobid...]</code>	When this job has begun execution, then jobs jobid... may begin.
<code>beforeok:jobid[:jobid...]</code>	If this job terminates execution without errors, then jobs jobid... may begin. See the csh warning under <a href="#">Extended description</a> .
<code>beforenotok:jobid[:jobid...]</code>	If this job terminates execution with errors, then jobs jobid... may begin. See the csh warning under <a href="#">Extended description</a> .
<code>beforeany:jobid[:jobid...]</code>	<p>When this job terminates execution, jobs jobid... may begin.</p> <p>If any of the before forms are used, the jobs referenced by jobid must have been submitted with a dependency type of on.</p> <p>If any of the before forms are used, the jobs referenced by jobid must have the same owner as the job being submitted. Otherwise, the dependency is ignored.</p>



Array dependencies make a job depend on an array or part of an array. If no count is given,

Dependency	Description
then the entire array is assumed. For examples, see <a href="#">Dependency examples</a> .	
<code>afterstartarray:arrayid[count]</code>	After this many jobs have started from arrayid, this job may start.
<code>afterokarray:arrayid[count]</code>	This job may be scheduled for execution only after jobs in arrayid have terminated with no errors.
<code>afternotokarray:arrayid[count]</code>	This job may be scheduled for execution only after jobs in arrayid have terminated with errors.
<code>afteranyarray:arrayid[count]</code>	This job may be scheduled for execution after jobs in arrayid have terminated, with or without errors.
<code>beforestartarray:arrayid[count]</code>	Before this many jobs have started from arrayid, this job may start.
<code>beforeokarray:arrayid[count]</code>	If this job terminates execution without errors, then jobs in arrayid may begin.
<code>beforenotokarray:arrayid[count]</code>	If this job terminates execution with errors, then jobs in arrayid may begin.
<code>beforeanyarray:arrayid[count]</code>	When this job terminates execution, jobs in arrayid may begin. If any of the before forms are used, the jobs referenced by arrayid must have been submitted with a dependency type of on. If any of the before forms are used, the jobs referenced by arrayid must have the same owner as the job being submitted. Otherwise, the dependency is ignored.
<p><b>i</b> Error processing of the existence, state, or condition of the job on which the newly submitted job is a deferred service, i.e. the check is performed after the job is queued. If an error is detected, the new job will be deleted by the server. Mail will be sent to the job submitter stating the error.</p>	

## Dependency examples

```
qsub -W depend=afterok:123.big.iron.com /tmp/script
```

```
qsub -W depend=before:234.hunk1.com:235.hunk1.com
```

```
/tmp/script
```

```
qsub script.sh -W depend=afterokarray:427[]
```

This assumes every job in array 427 has to finish successfully for the dependency to be satisfied.

```
qsub script.sh -W depend=afterokarray:427[][5]
```

This means that 5 of the jobs in array 427 have to successfully finish in order for the dependency to be satisfied.

---

## Operands

The `qsub` command accepts a script operand that is the path to the script of the job. If the path is relative, it will be expanded relative to the working directory of the `qsub` command.

If the script operand is not provided or the operand is the single character "-", the `qsub` command reads the script from standard input. When the script is being read from Standard Input, `qsub` will copy the file to a temporary file. This temporary file is passed to the library interface routine `pbs_submit`. The temporary file is removed by `qsub` after `pbs_submit` returns or upon the receipt of a signal which would cause `qsub` to terminate.

---

## Standard input

The `qsub` command reads the script for the job from standard input if the script operand is missing or is the single character "-".

---

## Input files

The script file is read by the `qsub` command. `qsub` acts upon any directives found in the script. When the job is created, a copy of the script file is made and that copy cannot be modified.

---

## Standard output

Unless the `-z` option is set, the job identifier assigned to the job will be written to standard output if the job is successfully created.

---

## Standard error

The `qsub` command will write a diagnostic message to standard error for each error occurrence.

---

## Environment variables

The values of some or all of the variables in the `qsub` commands environment are exported with the job (see the `-v` and `-V` options).

The environment variable `PBS_DEFAULT` defines the name of the default server. Typically, it

corresponds to the system name of the host on which the server is running. If PBS\_DEFAULT is not set, the default is defined by an administrator established file.

The environment variable PBS\_DPREFIX determines the prefix string which identifies directives in the script.

The environment variable PBS\_CLIENTRETRY defines the maximum number of seconds qsub will block (see the **-b** option). Despite the name, currently qsub is the only client that supports this option.

---

## torque.cfg

The `torque.cfg` file, located in PBS\_SERVER\_HOME (`/var/spool/torque` by default) controls the behavior of the `qsub` command. This file contains a list of parameters and values separated by whitespace.

- **QSUBSLEEP** – takes an integer operand which specifies time to sleep when running `qsub` command. Used to prevent users from overwhelming the scheduler.
- **SUBMITFILTER** – specifies the path to the submit filter used to pre-process job submission. The default path is `libexecdir/qsub_filter`, which falls back to `/usr/local/sbin/torque_submitfilter` for backwards compatibility. This `torque.cfg` parameter overrides this default.
- **SERVERHOST**
- **QSUBHOST**
- **QSUBSENDUID**
- **XAUTHPATH**
- **CLIENTRETRY**
- **VALIDATEGROUP**
- **DEFAULTCKPT**
- **VALIDATEPATH**
- **RERUNNABLEBYDEFAULT**

For example:

```
QSUBSLEEP 2
```

```
RERUNNABLEBYDEFAULT false
```

---

## Extended description

### Script Processing:

A job script may consist of PBS directives, comments and executable statements. A PBS directive provides a way of specifying job attributes in addition to the command line options. For example:

```
:
#PBS -N Job_name
#PBS -l walltime=10:30,mem=320kb
#PBS -m be
```

```
#
step1 arg1 arg2
step2 arg3 arg4
```

The `qsub` command scans the lines of the script file for directives. An initial line in the script that begins with the characters "`#!`" or the character ":" will be ignored and scanning will start with the next line. Scanning will continue until the first executable line, that is a line that is not blank, not a directive line, nor a line whose first non white space character is "#". If directives occur on subsequent lines, they will be ignored.

A line in the script file will be processed as a directive to `qsub` if and only if the string of characters starting with the first non white space character on the line and of the same length as the directive prefix matches the directive prefix.

The remainder of the directive line consists of the options to `qsub` in the same syntax as they appear on the command line. The option character is to be preceded with the "-" character.

If an option is present in both a directive and on the command line, that option and its argument, if any, will be ignored in the directive. The command line takes precedence.

If an option is present in a directive and not on the command line, that option and its argument, if any, will be processed as if it had occurred on the command line.

The directive prefix string will be determined in order of preference from:

- The value of the `-c` option argument if the option is specified on the command line.
- The value of the environment variable `PBS_DPREFIX` if it is defined.
- The four character string `#PBS`.

If the `-c` option is found in a directive in the script file, it will be ignored.

### User Authorization:

When the user submits a job from a system other than the one on which the PBS Server is running, the name under which the job is to be executed is selected according to the rules listed under the `-u` option. The user submitting the job must be authorized to run the job under the execution user name. This authorization is provided if:

- The host on which `qsub` is run is trusted by the execution host (see `/etc/hosts.equiv`).
- The execution user has an `.rhosts` file naming the submitting user on the submitting host.

### C-Shell .logout File:

The following warning applies for users of the c-shell, `csh`. If the job is executed under the `csh` and a `.logout` file exists in the home directory in which the job executes, the exit status of the job is that of the `.logout` script, not the job script. This may impact any inter-job dependencies. To preserve the job exit status, either remove the `.logout` file or place the following line as the first line in the `.logout` file:

```
set EXITVAL = $status
```

and the following line as the last executable line in `.logout`:

```
exit $EXITVAL
```

### Interactive Jobs:

If the `-I` option is specified on the command line or in a script directive, or if the "interactive" job

attribute declared true via the **-W** option, `-W interactive=true`, either on the command line or in a script directive, the job is an interactive job. The script will be processed for directives, but will not be included with the job. When the job begins execution, all input to the job is from the terminal session in which qsub is running.

When an interactive job is submitted, the qsub command will not terminate when the job is submitted. qsub will remain running until the job terminates, is aborted, or the user interrupts qsub with a SIGINT (the control-C key). If qsub is interrupted prior to job start, it will query if the user wishes to exit. If the user response "yes", qsub exits and the job is aborted.

Once the interactive job has started execution, input to and output from the job pass through qsub. Keyboard generated interrupts are passed to the job. Lines entered that begin with the tilde (~) character and contain special sequences are escaped by qsub. The recognized escape sequences are:

Sequence	Description
<code>~.</code>	qsub terminates execution. The batch job is also terminated.
<code>~susp</code>	Suspend the qsub program if running under the C shell. "susp" is the suspend character (usually Cntl-Z).
<code>~asusp</code>	Suspend the input half of qsub (terminal to job), but allow output to continue to be displayed. Only works under the C shell. "asusp" is the auxiliary suspend character, usually Cntl-Y.

---

## Exit status

Upon successful processing, the qsub exit status will be a value of zero.

If the qsub command fails, the command exits with a value greater than zero.

---

## Related topics

- [qalter\(1B\)](#)
- [qdel\(1B\)](#)
- [qhold\(1B\)](#)
- [qrls\(1B\)](#)
- [qsig\(1B\)](#)
- [qstat\(1B\)](#)
- [pbs\\_server\(8B\)](#)

## Non-Adaptive Computing topics

- [pbs\\_connect\(3B\)](#)
- [pbs\\_job\\_attributes\(7B\)](#)
- [pbs\\_queue\\_attributes\(7B\)](#)
- [pbs\\_resources\\_rix5\(7B\)](#)
- [pbs\\_resources\\_sp2\(7B\)](#)
- [pbs\\_resources\\_sunos4\(7B\)](#)
- [pbs\\_resources\\_unicos8\(7B\)](#)



pbs\_server\_attributes(7B)  
qselect(1B)  
qmove(1B)  
qmsg(1B)  
qrerun(1B)