

PBS Professional® 11.2

Programmer's Guide



PBS Works is a division of  Altair

PBS Professional Programmer's Guide; Altair PBS Professional 11.2, Updated: 12/16/11

Copyright © 2003-2011 Altair Engineering, Inc. All rights reserved.

PBS™, PBS Works™, PBS GridWorks®, PBS Professional®, PBS Analytics™, PBS Catalyst™, e-Compute™, and e-Render™ are trademarks of Altair Engineering, Inc. and are protected under U.S. and international laws and treaties. All other marks are the property of their respective owners.

ALTAIR ENGINEERING INC. Proprietary and Confidential. Contains Trade Secret Information. Not for use or disclosure outside ALTAIR and its licensed clients. Information contained herein shall not be decompiled, disassembled, duplicated or disclosed in whole or in part for any purpose. Usage of the software is only as explicitly permitted in the end user software license agreement.

Copyright notice does not imply publication.

For documentation and the PBS Works forums, go to:

Web: www.pbsworks.com

For more information, contact Altair at:

Email: pbssales@altair.com

Technical Support

Location	Telephone	e-mail
North America	+1 248 614 2425	pbssupport@altair.com
China	+86 (0)21 6117 1666	es@altair.com.cn
France	+33 (0)1 4133 0992	francesupport@altair.com
Germany	+49 (0)7031 6208 22	hwsupport@altair.de
India	+91 80 66 29 4500	pbs-support@india.altair.com
Italy	+39 0832 315573 +39 800 905595	support@altairengineering.it
Japan	+81 3 5396 2881	pbs@altairjp.co.jp
Korea	+82 31 728 8600	support@altair.co.kr
Scandinavia	+46 (0)46 286 2050	support@altair.se
UK	+44 (0)1926 468 600	pbssupport@uk.altair.com

This document is proprietary information of Altair Engineering, Inc.

Table of Contents

Acknowledgements	vii
About PBS Documentation	ix
1 Introduction	1
1.1 Location of API Libraries	1
1.2 Location of Header Files	1
1.3 Example Compilation Line	1
1.4 Deprecations	1
2 Concepts and Components	3
2.1 PBS Components	4
3 Server Functions	7
3.1 General Identifiers	7
3.2 Batch Server Functions	10
3.3 Server Management	10
3.4 Queue Management	13
3.5 Job Management	13
3.6 Server to Server Requests	19
3.7 Deferred Services	22
3.8 Resource Management	26
3.9 Network Protocol	27
4 Batch Interface Library (IFL)	29
4.1 Interface Library Overview	29
4.2 Interface Library Routines	30

Table of Contents

5	RPP Library	79
5.1	RPP Library Routines	79
6	TM Library	85
6.1	TM Library Routines.....	85
7	RM Library	93
7.1	RM Library Routines.....	93
8	TCL/tk Interface	97
8.1	TCL/tk API Functions.....	97
9	Hooks	105
9.1	Introduction	105
9.2	How Hooks Work	106
9.3	Interface to Hooks	107
9.4	The pbs Module	112
9.5	The pbs_stathook() API	121
10	HPC Basic Profile	123
10.1	Introduction	123
10.2	How PBS Works With HPC Basic Profile	123
10.3	Examples	128
10.4	Caveats	138
10.5	See Also	139
	Appendix A: License Agreement	141
	Index	151

List of Manual Pages

openrm, closerm, downrm, configrm, addreq, allreq, getreq, flushreq, activereq, fullresp	94
pbs_alterjob	31
pbs_connect	33
pbs_default	35
pbs_deljob	36
pbs_delresv	37
pbs_disconnect	38
pbs_geterrmsg	39
pbs_holdjob	40
pbs_locjob	41
pbs_manager	42
pbs_module	112
pbs_movejob	45
pbs_msgjob	46
pbs_orderjob	47
pbs_rerunjob	48
pbs_resreserve, pbs_resrelease	49
pbs_rlsjob	51
pbs_runjob, pbs_asyrunjob	52
pbs_selectjob	54
pbs_selstat	56
pbs_sigjob	59
pbs_stagein	60
pbs_stathook(3B)	121
pbs_statjob	61
pbs_statnode, pbs_statvnode, pbs_stathost	63
pbs_statque	65
pbs_statresv	67
pbs_statsched	69
pbs_statsserver	71
pbs_submit	73

pbs_submitresv	75
pbs_tclapi	98
pbs_terminate	77
rpp_open, rpp_bind, rpp_poll, rpp_io, rpp_read, rpp_write, rpp_close, rpp_getaddr, rpp_flush, rpp_terminate, rpp_shutdown, rpp_rcommit, rpp_wcommit, rpp_eom, rpp_getc, rpp_putc	80
tm_init, tm_nodeinfo, tm_poll, tm_notify, tm_spawn, tm_kill, tm_obit, tm_taskinfo, tm_atnode, tm_rescinfo, tm_publish, tm_subscribe, tm_finalize, tm_attach	86

Acknowledgements

PBS Professional is the enhanced commercial version of the PBS software originally developed for NASA. The NASA version had a number of corporate and individual contributors over the years, for which the PBS developers and PBS community are most grateful. Below we provide formal legal acknowledgements to corporate and government entities, then special thanks to individuals.

The NASA version of PBS contained software developed by NASA Ames Research Center, Lawrence Livermore National Laboratory, and MRJ Technology Solutions. In addition, it included software developed by the NetBSD Foundation, Inc., and its contributors, as well as software developed by the University of California, Berkeley and its contributors.

Other contributors to the NASA version of PBS include Bruce Kelly and Clark Streeter of NERSC; Kent Crispin and Terry Heidelberg of LLNL; John Kochmar and Rob Pennington of *Pittsburgh Supercomputing Center*; and Dirk Grunwald of *University of Colorado, Boulder*. The ports of PBS to the Cray T3e and the IBM SP SMP were funded by *DoD USAERDC*, Major Shared Research Center; the port of PBS to the Cray SV1 was funded by DoD MSIC.

No list of acknowledgements for PBS would possibly be complete without special recognition of the first two beta test sites. Thomas Milliman of the *Space Sciences Center* of the *University of New Hampshire* was the first beta tester. Wendy Lin of *Purdue University* was the second beta tester and holds the honor of submitting more problem reports than anyone else outside of NASA.



About PBS Documentation

Where to Keep the Documentation

To make cross-references work, put all of the PBS guides in the same directory.

What is PBS Professional?

PBS is a workload management system that provides a unified batch queuing and job management interface to a set of computing resources.

The PBS Professional Documentation

The documentation for PBS Professional includes the following:

PBS Professional Administrator's Guide:

Provides the PBS administrator with the information required to configure and manage PBS Professional (PBS).

PBS Professional Quick Start Guide:

Provides a quick overview of PBS Professional installation and license file generation.

PBS Professional Installation & Upgrade Guide:

Contains information on installing and upgrading PBS Professional.

PBS Professional User's Guide:

Covers user commands and how to submit, monitor, track, delete, and manipulate jobs.

PBS Professional Programmer's Guide:

Discusses the PBS application programming interface (API).

PBS Professional Reference Guide:

Contains PBS reference material.

PBS Manual Pages:

Describe PBS commands, resources, attributes, APIs

Ordering Software and Publications

To order additional copies of this manual and other PBS publications, or to purchase additional software licenses, contact your Altair sales representative. Contact information is included on the copyright page of this book.

Document Conventions

PBS documentation uses the following typographic conventions:

abbreviation

The shortest acceptable abbreviation of a command or subcommand is underlined.

command

Commands such as `qmgr` and `scp`

input

Command-line instructions

manpage (x)

File and path names. Manual page references include the section number in parentheses appended to the manual page name.

formats

Formats

Attributes

Attributes, parameters, objects, variable names, resources, types

Values

Keywords, instances, states, values, labels

Definitions

Terms being defined

Output

Output or example code

File contents



Chapter 1

Introduction

This book, the **Programmer's Guide** for PBS Professional, is provided to document the external application programming interfaces to the PBS Professional software.

1.1 Location of API Libraries

All of the libraries containing the PBS API are installed by default in `$PBS_EXEC/lib/`.

1.2 Location of Header Files

Header files used by customer-written code are found in `$PBS_EXEC/include`.

1.3 Example Compilation Line

An example of a compile command might look like the following:

```
cc mycode.c -I/usr/pbs/include -L/usr/pbs/lib -lpbs
```

1.4 Deprecations

The following are deprecated:

```
pbs_tclapi
```

```
pbs_resquery
```


Chapter 2

Concepts and Components

PBS is a distributed workload management system. As such, PBS handles the management and monitoring of the computational workload on a set of one or more computers. Modern workload/resource management solutions like PBS include the features of traditional batch queueing but offer greater flexibility and control than first generation batch systems (such as the original batch system NQS).

Workload management systems have three primary roles:

Queuing

The collecting together of work or tasks to be run on a computer. Users submit tasks or “jobs” to the resource management system where they are held until the system is ready to run them.

Scheduling

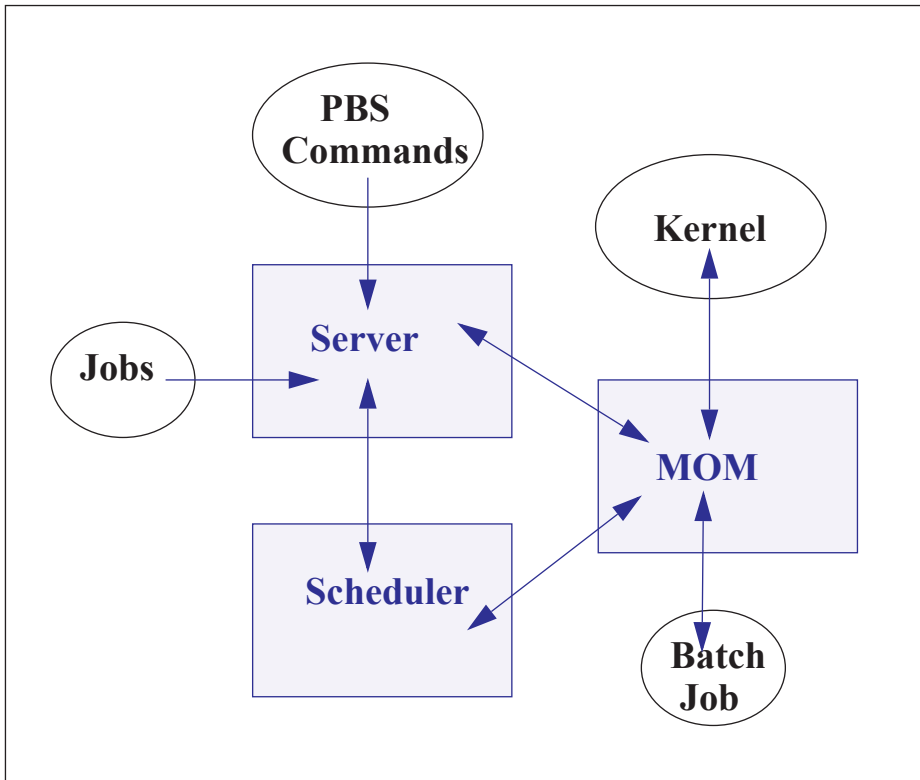
The process of selecting which jobs to run when and where, according to a predetermined policy. Sites balance competing needs and goals on the system(s) to maximize efficient use of resources (both computer time and people time).

Monitoring

The act of tracking and reserving system resources and enforcing usage policy. This covers both user-level and system-level monitoring as well as monitoring of the scheduling algorithms to see how well they are meeting the stated goals

2.1 PBS Components

PBS consist of two major component types: system daemons and user-level commands. A brief description of each is given here to help you make decisions during the installation process.



Job Server

The *Job Server* daemon process is the central focus for PBS. Within this document, it is generally referred to as *the Server* or by the execution name *pbs_server*. All commands and daemons communicate with the Server via an *Internet Protocol (IP)* network. The Server's main function is to provide the basic batch services such as receiving/creating a batch job, modifying the job, protecting the job against system crashes, and running the job. Typically there is one Server managing a given set of resources.

Job Executor (MOM)

The *Job Executor* is the daemon that actually places the job into execution. This daemon, *pbs_mom*, is informally called *MOM* as it is the mother of all executing jobs. (MOM is a reverse-engineered acronym that stands for Machine Oriented Miniserver.) MOM places a job into execution when it receives a copy of the job from a Server. MOM creates a new session that is as identical to a user login session as is possible. For example, if the user's login shell is `csh`, then MOM creates a session in which `.login` is run as well as `.cshrc`. MOM also has the responsibility for returning the job's output to the user when directed to do so by the Server. One MOM daemon runs on each computer which will execute PBS jobs.

A special version of MOM, called the *Globus MOM*, is available if it is enabled during the installation of PBS. It handles submission of jobs to the Globus environment. Globus is a software infrastructure that integrates geographically distributed computational and information resources. Globus is discussed in more detail in the “Globus Support” section of the **PBS Professional User's Guide**.

Job Scheduler

The *Job Scheduler* daemon, *pbs_sched*, implements the site's policy controlling when each job is run and on which resources. The Scheduler communicates with the various MOMs to query the state of system resources and with the Server to learn about the availability of jobs to execute. The interface to the Server is through the same API as used by the client commands. Note that the Scheduler communicates with the Server with the same privilege as the PBS Manager.

Commands

PBS supplies both command line programs that are POSIX 1003.2d conforming and a graphical interface. These are used to submit, monitor, modify, and delete jobs. These *client commands* can be installed on any system type supported by PBS and do not require the local presence of any of the other components of PBS.

There are three classifications of commands: user commands (which any authorized user can use), operator commands, and manager (or administrator) commands. Operator and Manager commands require specific access privileges, as discussed in the **PBS Professional Administrator's Guide**.

Chapter 3

Server Functions

This chapter presents formal definitions for identifiers and names to be used throughout the remainder of this document, followed by detailed discussion of the various functions of the PBS Professional Server process.

3.1 General Identifiers

The following identifiers or names are referenced throughout this document. Unless otherwise noted, their usage will conform to the definition and syntax described in the following subsections and to the general rules described in the next paragraph. If allowed as part of the identifier, when entering the identifier string on the command line or in a PBS job script directive, embedded single or double quote marks must be escaped by enclosing the string in the other type of quote mark. Therefore, the string may not contain both types of quote marks. If white space is allowed in the identifier string, the string must be quoted when it is entered on the command line or in a PBS job directive.

3.1.1 Account String

An Account String is a string of characters that some Server implementations may use to provide addition accounting or charge information. The syntax is unspecified except that it must be a single string. When provided on the command line to a PBS utility or in a directive in a PBS job script, any embedded white space must be escaped by enclosing the string in quotes.

3.1.2 Attribute Name

An Attribute Name identifies an attribute or data item that is part of the information that makes up a job, queue, or Server. The name must consist of alphanumeric characters plus the underscore, ‘_’, character. It should start with an alphanumeric character. The length is not limited. The names recognized by PBS are listed in sections 2.2, 2.3, and 2.4.

3.1.3 Destination Identifiers

A destination identifier is a string used to specify a particular destination. The identifier may be specified in one of three forms:

queue@server_name

queue

@server_name

where *queue* is an ASCII character string of up to 15 characters. Valid characters are alphanumerics, the hyphen and the underscore. The string must begin with a letter. *Queue* is the name of a queue at the batch Server specified by *server_name*. That Server will interpret the *queue* string. If *queue* is omitted, a null string is assumed. *server_name* is a string identifying a Server; see *server_name*, below. If *server_name* is omitted, the default Server is assumed.

3.1.4 Default Server

When a Server is not specified to a client, the client will send batch requests to the Server identified as the default Server. A client identifies the default Server by (a) the setting of the environment variable `PBS_DEFAULT` which contains a Server name, or (b) by editing the `PBS_SERVER` variable in the `/etc/pbs.conf` file on the local host. Note that if both are present, `PBS_DEFAULT` overrides the `PBS_SERVER` specification.

3.1.5 Host Name

A Host Name is a string that identifies a host or system on the network. The syntax of the string must follow the rules established by the network. For IP, a host name is of the form `name.domain`, where `domain` is a hierarchical, dot-separated List of subdomains. Therefore, a host name cannot contain a dot, “.” as a legal character other than as a subdomain separator. The name must not contain the commercial at sign, “@”, as this is often used to separate a file from the host in a remote file name. Also, to prevent confusion with port numbers (see section 2.7.9) a host name cannot contain a colon, “:”. The maximum length of a host name supported by PBS is defined by `PBS_MAXHOSTNAME`, currently set to 64.

3.1.6 Job Identifiers

When the term job identifier is used, the identifier is specified as:

`sequence_number[.server_name][@server]` The `sequence_number` is the number supplied by the Server when the job was submitted. The `server_name` component is the name of the Server which created the job. If it is missing, the name of the default Server will be assumed. `@server` specifies the current location of the job. When the term fully qualified job identifier is used, the identifier is specified as:

sequence_number.server[@server]

The `@server` suffix is not required if the job is still resides at the original Server which created the job. The `qsub` command will return a fully qualified job identifier.

3.1.7 Job Name

A Job Name is a string assigned by the user to provide a meaningful label to identify the job. The job name is up to and including 15 characters in length and may contain any printable characters other than white space. It must start with an alphanumeric character. If the user does not assign a name, PBS will assign a default name as described under the `-N` option of the `qsub(1)` command.

3.1.8 Resource Name

A Resource Name identifies a job resource requirement and may also identify a resource usage limit. The name must consist of alphanumeric characters plus the underscore, “_”, character. It should start with an alphanumeric character. The length is not limited. Certain resource names are identified and reserved by POSIX 1003.2d and by PBS. They are listed below in section “Types of Resources”.

3.1.9 Server Name.

Server Name is an ASCII character string of the form: `basic_server_name[:port]` The string identifies a batch Server. Basic Server names are identical to host names. The network routine `gethostbyname` will be used to translate to a network address. The network routine `getservbyname` will be used to determine the port number. An alternate port number may be specified by appending a colon, “:”, and the port number to the host name. This provides the means of specifying an alternate (test) Server on a host

3.1.10 User Name

A User Name is a string which identifies a user on the system under PBS. It is also known as the login name. PBS will accept names up to and including 16 characters. The name may contain any printable, non white space character excluding the commercial at sign, “@”. The various systems on which PBS is executing may place additional limitations on the user name.

3.2 Batch Server Functions

A batch Server provides services in one of two ways, (1) the Server provides a service at the request of a client; or (2) the Server provides a deferred service as a result of a change in conditions monitored by the Server. The Server also performs a number of internal bookkeeping functions that are described in this major section.

3.2.1 Client Service Requests

By definition, clients are processes that make requests of a batch Server. The requests may ask for an action to be performed on one or more jobs, one or more queues, or the Server itself. Those requests that cannot be successfully completed, are rejected. The reason for the rejection is returned in the reply to the client.

3.2.2 Deferred Services

The Server may, depending on conditions being monitored, defer a client service request until a later time. (Deferred services include file staging, job scheduling, etc.) Detailed discussion of the deferred services provided by the Server is given in section 3.7, “Deferred Services”, on page 22 below.

3.3 Server Management

The following sections describe the services provided by a batch Server in response to a request from a client. The requests are grouped in the following subsections by the type of object affected by the request: Server, queue, job, or resource. The batch requests described in

this section control the functioning of the batch Server. The control is either direct as in the Shut Down request, or indirect as when Server attributes are modified. The following table provides the numeric value of each of the batch request codes.

Table 3-1: Batch Request Codes

0	PBS_BATCH_Connect	24	PBS_BATCH_Rescq
1	PBS_BATCH_QueueJob	25	PBS_BATCH_ReserveResc
2	UNUSED	26	PBS_BATCH_ReleaseResc
3	PBS_BATCH_jobscript	27	PBS_BATCH_FailOver
4	PBS_BATCH_RdytoCommit	48	PBS_BATCH_StageIn
5	PBS_BATCH_Commit	49	PBS_BATCH_AuthenUser
6	PBS_BATCH_DeleteJob	50	PBS_BATCH_OrderJob
7	PBS_BATCH_HoldJob	51	PBS_BATCH_SelStat
8	PBS_BATCH_LocateJob	52	PBS_BATCH_RegistDep
9	PBS_BATCH_Manager	54	PBS_BATCH_CopyFiles
10	PBS_BATCH_MessJob	55	PBS_BATCH_DelFiles
11	PBS_BATCH_ModifyJob	56	PBS_BATCH_JobObit
12	PBS_BATCH_MoveJob	57	PBS_BATCH_MvJobFile
13	PBS_BATCH_ReleaseJob	58	PBS_BATCH_StatusNode
14	PBS_BATCH_Rerun	59	PBS_BATCH_Disconnect
15	PBS_BATCH_RunJob	60-61	UNUSED
16	PBS_BATCH_SelectJobs	62	PBS_BATCH_JobCred
17	PBS_BATCH_Shutdown	63	PBS_BATCH_CopyFiles_Cred
18	PBS_BATCH_SignalJob	64	PBS_BATCH_DelFiles_Cred
19	PBS_BATCH_StatusJob	65	PBS_BATCH_GSS_Context
20	PBS_BATCH_StatusQue	66-69	UNUSED
21	PBS_BATCH_StatusSvr	70	PBS_BATCH_SubmitResv

Table 3-1: Batch Request Codes

22	PBS_BATCH_TrackJob	71	PBS_BATCH_StatusResv
23	PBS_BATCH_AsyrunJob	72	PBS_BATCH_DeleteResv

3.3.1 Manage Request

The Manage request supports the `qmgr (8)` command and several of the operator commands. The command directs the Server to create, alter, or delete an object managed by the Server or one of its attributes. For more information, see the `qmgr` command.

3.3.2 Server Status Request

The status of the Server may be requested with a Server Status request. The batch Server will reject the request if the user of the client is not authorized to query the status of the Server. If the request is accepted, the Server will return a Server Status Reply. See the `qstat` command and the Data Exchange Format description for details of which Server attributes are returned to the client.

3.3.3 Start Up

A batch request to start a Server cannot be sent to a Server since the Server is not running. Therefore a batch Server must be started by a process local to the host on which the Server is to run. The Server is started by a `pbs_server` command. The Server recovers the state of managed objects, such as queues and jobs, from the information last recorded by the Server. The treatment of jobs which were in the running state when the Server previously shut down is dictated by the start up mode, see the description of the `pbs_server (8)` command.

3.3.4 Shut Down

The batch Server is "shut down" when it no longer responds to requests from clients and does not perform deferred services. The batch Server is requested to shut down by sending it a Server Shutdown request. The Server will reject the request from a client not authorized to shut down the Server. When the Server accepts a shut down request, it will terminate in the manner described under the `qterm` command. When shutting down, the Server must record the state of all managed objects (jobs, queues, etc.) in non-volatile memory. Jobs which were running will be marked in the secondary state field for possible special treatment when the Server is restarted. If checkpoint is supported, any job running at the time of the shut down request whose Checkpoint attribute is not n, will be checkpointed. This includes jobs whose

Checkpoint attribute value is “unspecified”, a value of u. If the Server receives either a SIGTERM or a SIGSHUTDOWN signal, the Server will act as if it had received a shut down immediate request.

3.4 Queue Management

The following client requests effect one or more queues managed by the Server. These requests require a privilege level generally assigned to operators and administrators.

3.4.1 Queue Status Request

The status of a queue at the Server may be requested with a Queue Status request. The batch Server will reject the request if any of the following conditions are true:

- The user of the client is not authorized to query the status of the designated queue.
- The designated queue does not exist on the Server.

If the request does not specify a queue, status of all the queues at the Server will be returned. When the request is accepted, the Server will return a Queue Status Reply. See the `qstat` command and the Data Exchange Format description for details of which queue attributes are returned to the client.

3.5 Job Management

The following client requests effect one or more jobs managed by the Server. These requests do not require any special privilege except when the job for which the request is issued is not owned by the user making the request.

3.5.1 Queue Job Request

A Queue Job request is a complex request consisting of several subrequests: Initiate Job Transfer, Job Data, Job Script, and Commit. The end result of a successful Queue Job request is an additional job being managed by the Server. The job may have been created by the request or it may have been moved from another Server. The job resides in a queue managed by the Server. When a queue is not specified in the request, the job is placed in a queue selected by the Server. This queue is known as the default queue. The default queue is an

attribute of the Server that is settable by the administrator. The queue, whether specified or defaulted, is called the target queue. The batch Server will reject a Queue Job Request if any of the following conditions are true:

- The client is not authorized to create a job in the target queue.
- The target queue does not exist at the Server.
- The target queue is not enabled.
- The target queue is an execution queue and a resource requirement of the job exceeds the limits set upon the queue.
- The target queue is an execution queue and an unrecognized resource is requested by the job.
- The job requires access to a user identifier that the client is not authorized to access.

When a job is placed in a execution queue, it is placed in the queued state unless one of the following conditions applies:

- The job has an `execution_time` attribute that specifies a time in the future and the `Hold_Types` attribute has value of `{NONE}`; in which case the job is placed in the waiting state.
- The job has a `Hold_Types` attribute with a value other than `{NONE}`, wherein the job is placed in the held state.

When a job is placed in a routing queue, its state may change based on the conditions described in section 3.7.4, “Job Routing”, on page 24.

A Server that accepts a Queue Job Request for a new job will: (1) add the `PBS_O_QUEUE` variable to the `Variable_List` attribute of the job and set the value to the name of the target queue; (2) add the `PBS_JOBID` variable to the `Variable_List` attribute of the job and set the value to the job identifier assigned to the job; (3) add the `PBS_JOBNAME` variable to the `Variable_List` attribute of the job and set the value to the value of the `Job_Name` attribute of the job. When the Server accepts a Queue Job request for an existing job, the Server will send a Track Job request to the Server which created the job.

3.5.2 Job Credential Request

The Job Credential sub-request is part of the Queue Job complex request. This sub-request transfers a copy of the credential provided by the authentication facility explained below.

3.5.3 Job Script Request

The Job Script sub-request is part of the Queue Job complex request. This sub-request passes a block of the job script file to the receiving Server. The script is broken into 8 kilobyte blocks to prevent having to hold the entire script in memory. One or more Job Script sub-requests may be required to transfer the script file.

3.5.4 Commit Request

The Commit sub-request is part of the Queue Job request. The Commit notifies the receiving Server that all parts of the job have been transferred and the receiving Server should now assume ownership of the job. Prior to sending the Commit, the sending client, command or another Server, is the owner.

3.5.5 Message Job Request

A batch Server can be requested to write a string of characters to one or both output streams of an executing job. This request is primarily used by an operator to record a message for the user. The batch Server will reject a Message Job request if any of the following conditions are true:

- The designated job is not in the running state.
- The user of the client is not authorized to post a message to the designated job.
- The designated job is not owned by the Server.

When the Server accepts the Message Job request, it will forward the request to the primary MOM daemon for the job. (Upon receipt of the Message Job request from the Server, the MOM will append the message string, followed by a new line character, to the file or files indicated. If no file is indicated, the message will be written to the standard error of the job.)

3.5.6 Locate Job Request

A client may ask a Server to respond with the location of a job that was created or is owned by the Server. When the Server accepts the Locate Job request, it returns a Locate Reply. The request will be rejected if any of the following conditions are true:

- The Server does not own (manage) the job, and
- The Server did not create the job.
- The Server is not maintaining a record of the current location of the job.

3.5.7 Delete Job Request

A Delete Job request asks a Server to remove a job from the queue in which it exists and not place it elsewhere. The batch Server will reject a Delete Job Request if any of the following conditions are true:

- The user of the client is not authorized to delete the designated job.
- The designated job is not owned by the Server.
- The designated job is not in an eligible state. Eligible states are queued, held, waiting, running, and transiting.

If the job is in the running state, the Server will forward the Delete Job request to the primary MOM daemon responsible for the job. (Upon receipt, the MOM daemon will first send a SIGTERM signal to the job process group. After a delay specified by the delete request or if not specified, the `kill_delay` queue attribute, the MOM will send a SIGKILL signal to the job process group. The job is then placed into the exiting state.) Option arguments exist to specify the “delay” time (seconds) between the SIGTERM and SIGKILL signals, as well as to “force” the deletion of the job even if the node on it is running is not responding.

3.5.8 Modify Job Request

A batch client makes a Modify Job request to the Server to alter the attributes of a job. The batch Server will reject a Modify Job Request if any of the following conditions are true:

- The user of the client is not authorized to make the requested modification to the job.
- The designated job is not owned by the Server.
- The requested modification is inconsistent with the state of the job.
- A requested resource change would exceed the limits of the queue or Server.
- An unrecognized resource is requested for a job in an execution queue.

When the batch Server accepts a Modify Job Request, it will modify all the specified attributes of the job. When the batch Server rejects a Modify Job Request, it will modify none of the attributes of the job.

3.5.9 Run Job

The "Run Job" request directs the Server to place the specified job into immediate execution. The request is issued by a `qrun` operator command and by the PBS Job Scheduler.

3.5.9.1 Rerun Job Request

To rerun a job is to kill the members of the session (process) group of the job and leave the job in the execution queue. If the `Hold_Types` attribute is not `{NONE}`, the job is eligible to be re-scheduled for execution. The Server will reject the Rerun Job request if any of the following conditions are true:

- The user of the client is not authorized to rerun the designated job.
- The `Rerunnable` attribute of the job has the value `{FALSE}`.
- The job is not in the running state.
- The Server does not own the job.

When the Server accepts the Rerun Job request, the request will be forwarded to the primary MOM responsible for the job, who will then perform the following actions:

- Send a `SIGKILL` signal to the session (process) group of the job.
- Send an `OBIT` notice to the Server with resource usage information
- The Server will then requeue the job in the execution queue in which it was executing.

If the `Hold_Types` attribute is not `{NONE}`, the job will be placed in the held state. If the `execution_time` attribute is a future time, the job will be placed in the waiting state. Otherwise, the job is placed in the queued state.

3.5.10 Hold Job Request

A client can request that one or more holds be applied to a job. The batch Server will reject a Hold Job request if any of the following conditions are true:

- The user of the client is not authorized to add any of the specified holds.
- The batch Server does not manage the specified job.

When the Server accepts the Hold Job Request, it will add each type of hold listed which is not already present to the value of the `Hold_Types` attribute of the job. If the job is in the queued or waiting state, it is placed in the held state. If the job is in running state, then the following additional actions are taken: If check-point / restart is supported by the host system, placing a hold on a running job will cause the job (1) to be checkpointed, (2) the resources assigned to the job will be released, and (3) the job is placed in the held state in the execution queue. If checkpoint / restart is not supported, the Server will only set the requested hold attribute. This will have no effect unless the job is rerun or restarted.

3.5.11 Release Job Request

A client can request that one or more holds be removed from a job. A batch Server rejects a Release Job request if any of the following conditions are true:

- The user of the client is not authorized to add (remove) any of the specified holds.
- The batch Server does not manage the specified job.

When the Server accepts the Release Job Request, it will remove each type of hold listed from the value of the `Hold_Types` attribute of the job. Normally, the job will then be placed in the queued state, unless another hold type is remaining on the job. However, if the job is in the held state and all holds have been removed, the job is placed in the waiting state if the `Execution_Time` attribute specifies a time in the future.

3.5.12 Move Job Request

A client can request a Server to move a job to a new destination. The batch Server will reject a Move Job Request if any of the following conditions are true:

- The user of the client is not authorized to remove the designated job from the queue in which the job resides.
- The user of the client is not authorized to submit a job to the new destination.
- The designated job is not owned by the Server.
- The designated job is not in the queued, held, or waiting state.
- The new destination is disabled.
- The new destination is inaccessible. When the Server accepts a Move Job request, it will
 - Queue the designated job at the new destination.
 - Remove the job from the current queue.

If the destination exists at a different Server, the current Server will transfer the job to the new Server by sending a Queue Job request sequence to the target Server. The Server will insure that a job is neither lost nor duplicated.

3.5.13 Select Jobs Request

A client is able to request from the Server a list of jobs owned by that Server that match a list of selection criteria. The request is a Select Jobs request. All the jobs owned by the Server and which the user is authorized to query are initially eligible for selection. Job attributes and resources relationships listed in the request restrict the selection of jobs. Only jobs which have attributes and resources that meet the specified relations will be selected. The Server will

reject the request if the queue portion of a specified destination does not exist on the Server. When the request is accepted, the Server will return a Select Reply containing a list of zero or more jobs that met the selection criteria.

3.5.14 Signal Job Request

A batch client is able to request that the Server signal the session (process) group of a job. Such a request is called a Signal Job request. The batch Server will reject a Signal Job Request if any of the following conditions are true:

- The user of the client is not authorized to signal the job.
- The job is not in the running state, except for the special signal “resume” when the job must be in the Suspended state.
- The Server does not own the designated job.
- The requested signal is not supported by the host operating system. (The kill system call returns [EINVAL].)

When the Server accepts a request to signal a job, it will forward the request to the primary MOM daemon responsible for the job, who will then send the signal requested by the client to the all processes in the job’s session.

3.5.15 Status Job Request

The status of a job or set of jobs at a destination may be requested with a Status Job request. The batch Server will reject a Status Job Request if any of the following conditions are true:

- The user of the client is not authorized to query the status of the designated job.
- The designated job is not owned by the Server.

When the Server accepts the request, it will return a Job Status Message to the client. See the `qstat` command and the Data Exchange Format description for details of which job attributes are returned to the client. If the request specifies a job identifier, status will be returned only for that job. If the request specifies a destination identifier, status will be returned for all jobs residing within the specified queue that the user is authorized to query.

3.6 Server to Server Requests

Server to Server requests are a special category of client requests. They are only issued to a Server by another Server.

3.6.1 Track Job Request

A client that wishes to request an action be performed on a job must send a batch request to the Server that currently manages the job. As jobs are routed or moved through the batch network, finding the location of the job can be difficult without a tracking service. The Track Job request forms the basis for this service. A Server that queues a job sends a track job request to the Server which created the job. Additional backup location Servers may be defined. A Server that receives a track job request records the information contained therein. This information is made available in response to a Locate Job request.

3.6.2 Synchronize Job Starts

PBS provides for synchronizing the initiation of separate jobs. This is done to support distributing processing. Job start synchronization is requested through a special dependency attribute. The first job in the set, the “master”, specifies the dependency attribute as:

```
-W synccount=count
```

where `count` is an integer which is the number of other jobs to be synchronized with this job. This job is the master only in the sense that it defines the rendezvous point for the semaphore messages and that it must be submitted first so the identifier is known for the other jobs in the set. The other jobs in the sync set specify the dependency attribute as:

```
-W syncwith=job_identifier
```

where `job_identifier` is the job identifier assigned to the job which contained the sync-count resource, the master job. When the Server queues a job in an execution queue and the job is a member of a sync set, including the “master”, the Server places a system hold on the job. The secondary state is set to indicate the system hold is for sync. The Server managing the non master jobs will register the job with the Server managing the master by sending a Register Dependent request with a "Register" operation. When all jobs have registered, as determined by the count on the master, the Server managing the master job will send a Register Dependent request, with a "Release" operation, request to each job in turn in the set to remove the system hold. The released jobs may now vie for resources. The jobs are released in order of the “cheapest” resources first; the concept of “Resource Costs” will be explained shortly. When the resources required by a released job are available, as determined by the Scheduler, A run Job Request will be issued for that job. The Server which manages the job will send a Register Dependent request with a “Ready” operation to the Server that owns the master job. This request indicates that the dependent job is ready and the job with the next cheapest resources can be released.

If the master of a sync set is aborted before all jobs in the set begin execution, an Abort Job request is sent to all jobs in the set. This is done because the synchronous feature is intended for a set jobs which need communication amongst themselves during execution. If the master is gone, (1) the rendezvous point for Server messages is lost, and (2) the job set is unlikely to be able to establish the inter job communications required.

3.6.3 Job Dependency

PBS provides support for job dependency. A job, the “child”, can be declared to be dependent on one or more jobs, the “parents”. A parent may have any number of children. The dependency is specified as an attribute on the `qsub` command with the `-W` option. The general specification is of the form:

```
-W type=argument[, type=argument, ...]
```

See the `qalter(1B)` or `qsub(1B)` man pages for the complete specification of the dependency list, and the **PBS Professional User’s Guide** for detailed discussion of use.

When a Server queues a job with a dependency type of `syncwith`, `after`, `afterok`, `after notok`, or `after-any` in an execution queue, the Server will send a Register Dependent Job request to the Server managing the job specified by the associated `job_identifier`. The request will specify that the Server is to register the dependency. This actually creates a corresponding `before` type dependency attribute entry on the parent (e.g. run job X *before* job Y). If the request is rejected because the parent job does not exist, the child job is aborted. If the request is accepted, a system hold is placed on the child job. When a parent job, with any of the `before . . .` types of dependency, reaches the required state, started or terminated, the Server executing the parent job sends a Register Dependent Job request to the Server managing the child job directing it to release the child job. If there are no other dependencies on other jobs, the system hold on the child job is removed. When a child job is submitted with an `on` dependency and the parent is submitted with any of the `before . . .` types of dependencies, the parent will register with the child. This causes the `on` dependency count to be reduced and a corresponding `after . . .` dependency to be created for the child job. The result is a pairing between corresponding `before . . .` and `after . . .` dependency types. If the parent job terminates in a manner that the child is not released, it is up to the user to correct the situation by either deleting the child job or by correcting the problem with the parent job and resubmitting it. If the parent job is resubmitted, it must have a dependency type of `before`, `beforeok`, `beforenotok`, or `beforeany` specified to connect it to the waiting child job.

3.7 Deferred Services

This section describes the deferred services performed by batch Servers: file staging, job selection, job initiation, job routing, job exit, job abort, and the rerunning of jobs after a restart of the Server. The following rules apply to deferred services on behalf of jobs:

- If the Server cannot complete a deferred service for a reason which is permanent, then the job is aborted.
- If the service cannot be completed at the current time but may be later, the service is retried a finite number of times.

3.7.1 Job Scheduling

If the Server attribute `scheduling` is set true, the Server will immediately request a scheduling cycle of the PBS Job Scheduler. While it remains true, the Scheduler will be cycled when any of four events occur:

- Enqueuing of a job in an execution queue or the change of state of a job in an execution queue to Queued from Waiting or Held.
- Termination of a running job. The termination may be normal execution completion, or because the job was deleted by request.
- Elapse of a specified cycle time as established by the administrator.
- The completion of a scheduling cycle in which one and only one job was scheduled for execution. This provides for the implementation of scheduling scripts that must see the impact of the new job on system resources before picking a second job.

While a request for a scheduling cycle is outstanding, the connection to the Scheduler is open, the Server will not make another request of the Scheduler. If the Server attribute `scheduling` is set false, the Server will not contact the scheduler. This condition is indicated by the `server_state` attribute as Idle.

3.7.2 File Staging

Two types of file staging services exist, in-staging before execution and out-staging after execution. These services are requested by an attribute (via the `-w` option) which specifies the files to be staged:

```
-wstagein=local_file@host:remote_path [,local_file@host:remote_path,...]  
-wstageout=local_file@host:remote_path [,local_file@host:remote_path,...]
```

A request to stage in a file directs the Server to direct MOM to copy a file from a remote host to the local host. The user must have authority to access the file under the same user name under which the job will be run. The remote file is not modified or destroyed. The file will be available before the job is initiated. If a file cannot be staged in for any reason, any files which were staged-in are deleted and the job is placed into wait state and mail is sent to the job owner.

A request to stage out a file directs the Server to direct MOM to move a file from the local host to a remote host. This service is performed after the job has completed execution and regardless of its exit status. If a file cannot be moved, mail is sent to the job owner. If a file is successfully staged out, the local file is deleted. A version of the BSD 4.4-Lite system utility, `rcp(1)`, will be used to move files over the network. This version of `rcp` has been modified to always return a non-zero exit status on any failure.

3.7.3 Job Initiation

Job initiation is to place a job into execution. The Server may receive a Run Job request from the `qrun` command, or the PBS Job Scheduler. If the request is authenticated, then the Server forwards the Run Job request to the appropriate MOM (as either specified in the Run Job request, or as selected by the Server itself if unspecified).

The receiving MOM daemon will then create a session leader that runs the shell program indicated by the `Shell_Path_List` attribute of the job. The pathname of the script and any script arguments are passed as parameters to the shell. If the path name of the shell is a relative name, the MOM will search its execution path, `$PATH`, for the shell. If the path name of the shell is omitted or is the null string, the MOM uses the login shell for the user under whose name the job is to be run. The MOM will determine the user name under which the job is to be run by the following rules:

1. Select the user identifier from the `User_List` job attribute which has a host name that matches the execution host.
2. Select the user identifier from the `User_List` job attribute which has no associated host name.
3. Use the user name from the `job_owner` attribute of the job.

The MOM will create, in the environment of the session leader of the job, the environment variables named: `PBS_ENVIRONMENT`, the value of which is the string “`PBS_BATCH`”. `PBS_QUEUE` has the value of the name of the execution queue. The MOM will also place in the environment of the session leader of the job, all of the variables and their corresponding values found in the `variables` attribute of the job. The MOM will place the required limits on the resources for which the host system supports resource limits. If the job had been run before and is now being rerun, the MOM will insure that the standard output and standard error streams of the job are appended to the prior streams, if any. If the MOM and host system

support accounting, the MOM will use the value of the `Account_Name` job attribute as required by the host system. If the MOM and host system support checkpoint, the MOM will set up checkpointing of the job according to the value of the `Checkpoint` job attribute. If checkpoint is supported and the `Checkpoint` attribute requests checkpointing at the minimum interval or a interval less than the minimum interval for the queue, then checkpoint will be set for an interval given by the queue attribute `minimum_interval`. The MOM will set up the standard output stream and the standard error stream of the job according to the following rules:

- The stream will be located in a temporary file in the MOM's `spool` directory.
- If the job attribute `Join_Path` has the value `eo` or the value `oe`, the MOM connects the standard error stream of the job to the same file as the standard output stream.

3.7.4 Job Routing

Job routing is moving a job from a routing queue to one of the destinations associated with the queue. If the `started` queue attribute is `{TRUE}`, the Server will route all eligible jobs which reside in the queue. All jobs in the `queued` state are eligible. If the queue attribute `route_held_jobs` is `{TRUE}`, jobs in the `held` state are eligible for routing. If the queue attribute `route_waiting_jobs` is `{TRUE}`, jobs in the `waiting` state are eligible. The Server will execute the function specified by the queue attribute `route_function` to select a destination for the job. Possible destinations are listed in the queue attribute `route_destinations`. If the destination to which the job is to be routed is at another Server, the current Server will use a `Queue Job` request sequence to move the job to the new destination. If the Server is unable to route a job to a chosen destination, the Server will select another destination from the list and retry the route. If the Server is unable to route a job to any destination because of a temporary condition, such as being unable to connect with the Server at the destination, the Server will retry the route after a delay specified by the queue attribute `route_retry_time`. The Server will proceed to route other jobs in the queue. The Server will retry the route up to the (queue attribute) `number_retries` times. If the Server is unable to route a job to any destination and all failures are permanent (non-temporary), the Server will abort the job.

3.7.5 Job Exit

When the session leader of a batch job exits, the MOM will perform the following actions in the order listed.

- Place the job in the `exiting` state.
- “Free” the resources allocated to the job. The actual releasing of resources assigned to the processes of the job is performed by the kernel. PBS will free the resources which it

“reserved” for the job by decrementing the `resources_used` generic data item for the queue and Server.

- Return the standard output and standard error streams of the job to the user. If the `Keep_Files` attribute of the job contains `{KEEP_OUTPUT}`, the Server copies the spooled file holding the standard output stream of the job to the home directory of the user under whose name the job executed. The file name for the output is `job_name.oseq_number`. See the `qsub(1B)` command description. If the `Keep_Files` attribute of the job contains `{KEEP_ERROR}` and the `Join_Path` attribute does not contain `'e'`, the Server copies the spooled file holding the standard error stream of the job to the home directory of the user under whose name the job executed. The file name for the error file is `job_name.eseq_number`.

If the files are not to be kept on the execution host as described above, the temporary file holding the standard output is copied or renamed to the host and path name specified by the job attribute `Output_Path`. If the path name is relative, the file will be located relative to home directory of the user on the receiving host.

- If the `Join_Path` attribute does not contain the value `e`, the standard error of the job is delivered according to the same rules as the standard output described above. If either output file cannot be copied to its specified destination, the Server will send mail to the job owner specifying the current location of the output.
- If the `Mail_Points` job attribute contains the value `{EXIT}`, the Server will send mail to the users listed in the job attribute `Mail_List`.
- If out staging of files is supported, the files listed in the outfile resource will be copied to the specified destination.
- The job will be removed from the execution queue.

3.7.6 Job Aborts

If the Server aborts a job and the `Mail_Points` job attribute contains the value `{ABORT}`, the Server will send mail to the users listed in the job attribute `Mail_List`. The mail message will contain the reason the job was aborted. In addition, the `stdout` and `stderr` files specified for the job, if they exist, will be copied back to the specified location.

3.7.7 Timed Events

The Server performs certain events at a specified time or after a specified time delay. A job may have an `execution_time` attribute set to a time in the future. When that time is reached, the job state is updated. If the Server is unable to make connection with another Server, it is to retry after a time specified by the routing queue attribute `route_retry_time`.

3.7.8 Event Logging

The PBS Server maintains an event logfile, the format and contents of which are documented in the **PBS Professional Administrator's Guide**.

3.7.9 Accounting.

The PBS Server maintains an accounting file, the format and contents of which are documented in the **PBS Professional Administrator's Guide**.

3.8 Resource Management

PBS performs resource allocation at job initiation in two ways depending on the support provided by the host system. Resources are either reservable or non reservable.

3.8.1 Resource Limits

When submitting a job, a user may specify the hard limit of usage for resources known to the system on which the job will run. If the executing job usage of resources exceed the specified limit, the job is aborted. If the user does not specify a limit for a resource type, the limit may be set to a default established by the PBS administrator. The default limit is taken from the first of the following attributes which is set:

1. The current queue's attribute `resources_default`.
2. The Server's attribute `resources_default`.
3. The current queue's attribute `resources_max`.
4. The Server's attribute `resources_max`.

If the user does not specify a limit for a resource and a default is not established via one of the above attributes, the usage of the resource is unlimited.

3.8.2 Resource Names

For additional information, see the **PBS Professional User's Guide** where all resource names are documented.

3.9 Network Protocol

The PBS system fits into a client - Server model, with a batch client making a request of a batch Server and the Server replying. This client - Server communication necessitates an interprocess communication method and a data exchange (data encoding) format. Since the client and Server may reside on different systems, the interprocess communication must be supportable over a network.

While the basic PBS system fits nicely into the client - Server model, it also has aspects of a transaction system. When jobs are being moved between Servers, it is critical that the jobs are not lost or replicated. Updates to a batch job must be applied once and only once. Thus the operation must be atomic. Most of the client to Server requests consist of a single message. Treating these requests as an atomic operation is simple. One request, "Queue Job", is more complex and involves several messages, or subrequests, between the client and the Server. Any of these subrequests might be rejected by the Server. It is important that either side of the connection be able to abort the request (transaction) without losing or replicating the job. The network connection also might be lost during the request. Recovery from a partially transmitted request sequence is critical. The sequence of recovery from lost connections is discussed in the Queue Job Request description.

The batch system data exchange protocol must be built on top of a reliable stream connection protocol. PBS uses TCP/IP and the socket interface to the network. Either the Simple Network Interface, SNI, or the Detailed Network Interface, DNI, as specified by POSIX.12, Protocol Independent Interfaces, could be used as a replacement.

3.9.1 General DIS Data Encoding

The purpose of the "Data is Strings" encoding is to provide a simple, fast, small, machine independent form for encoding data to a character string and back again. Because data can be decoded directly into the final internal data structures, the number of data copy operations are reduced. Data items are represented as people think of them, but preceded with a count of the length of each data item. For small positive integers, it is impossible to tell from the encoded data whether they came from `signed chars`, `shorts`, `ints`, or `longs`. Similarly, for small negative numbers, the only thing that can be determined from the encoded data is that the source datum was not unsigned. It is impossible to tell the word size of the encoding machine, or whether it uses 2's complement, one's complement or sign - magnitude representation, or even if it uses binary arithmetic. All of the basic C data types are handled. Signed and unsigned chars, shorts, ints, longs produce integers. NULL terminated and counted strings produce counted strings (with the terminating NULL removed). Floats, doubles, and long doubles produce real numbers. Complex data must be built up from the basic types. Note that there is no type tagging, so the type and sequence of data to be decoded must be known in advance.

Chapter 4

Batch Interface Library (IFL)

The primary external application programming interface to PBS is the Batch Interface Library, or IFL. This library provides a means of building new batch clients. Any batch service request can be invoked through calls to the batch interface library. Users may wish to build a job which could status itself or spawn off new jobs. Or they may wish to customize the job status display rather than use `qstat`. Administrators may use the interface library to build new control commands.

4.1 Interface Library Overview

The IFL provides a user-callable function corresponding to each batch client command. There is (approximately) a one to one correlation between commands and batch service requests. Additional routines are provided for network connection management. The user callable routines are declared in the header file `PBS_ifl.h`. Users open a connection with a batch Server via a call to `pbs_connect()`. Multiple connections are supported. Before a connection is established, `pbs_connect()` will fork and exec an `pbs_iff` process, as shown in figure 4-1 below. The purpose of `pbs_iff` is to provide the user a credential which validates the user's identity. This credential is included in each batch request. The provided credential prevents a user from spoofing another user's identity.

The credential that is sent to the server consists of: a) user's name from the password file based on running `pbs_iff`'s "real uid" value, and b) unprivileged, client-side port value associated with the original `pbs_connect` request message to the server. The server looks at the entries in its connection table to try and find the entry having these two pieces of information, and which is not yet marked authenticated. To be believed, this information must be gotten from a connection having a privileged, remote-end, port value.

After all requests have been made to a Server, its connection is closed via a call to `pbs_disconnect()`.




Figure 4-1: Interface Between Client, IFF, and Server

Users request service of a batch Server by calling the appropriate library routine and passing it the required parameters. The parameters correspond to the options and operands on the commands. It is the user's responsibility to ensure the parameters have correct syntax. Each function will return zero upon success and a non-zero error code on failure. These error codes are available in the header file `PBS_error.h`. The library routine will accept the parameters and build the corresponding batch request, then pass it to the Server.

To use `pbs_connect` with Windows, initialize the network library and link with `winsock2`. Call `winsock_init()` before calling `pbs_connect()`, and link against the `ws2_32.lib` library.

Any user-written programs using the IFL API must link with the `pthread` library.

4.2 Interface Library Routines

The following manual pages describe the user-callable functions in the IFL.

pbs_alterjob

alter pbs batch job

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_alterjob(int connect, char *job_id, struct attrl *attrib,
char *extend)
```

DESCRIPTION

Issue a batch request to alter a batch job.

A Modify Job batch request is generated and sent to the server over the connection specified by connect which is the return value of pbs_connect().

The argument, job_id , identifies which job is to be altered. It is specified in the form:

```
sequence_number.server
```

The parameter, attrib , is a pointer to an attrl structure which is defined in pbs_ifl.h as:

```
struct attrl {
    char *name;
    char *resource;
    char *value
    struct attrl *next;
};
```

The attrib list is terminated by the first entry where next is a null pointer.

The name member points to a string which is the name of the attribute. The value member points to a string which is the value of the attribute. The attribute names are defined in pbs_ifl.h.

If attrib itself is a null pointer, then no attributes are altered.

Associated with an attribute of type `ATTR_1` (the letter `ell`) is a resource name indicated by `resource` in the `attr1` structure. All other attribute types should have a pointer to a null string ("") for `resource` .

If the `resource` of the specified resource name is already present in the job's `Resource_List` attribute, it will be altered to the specified value. If the `resource` is not present in the attribute, it is added.

Certain attributes of a job may or may not be alterable depending on the state of the job; see `qalter(1B)`.

The parameter, `extend` , is reserved for implementation defined extensions.

SEE ALSO

`qalter(1B)`, `qhold(1B)`, `qrls(1B)`, `qsub(1B)`, `pbs_connect(3B)`, `pbs_holdjob(3B)`, and `pbs_rlsjob(3B)`

DIAGNOSTICS

When the batch request generated by `pbs_alterjob()` function has been completed successfully by a batch server, the routine will return 0 (zero). Otherwise, a non zero error is returned. The error number is also set in `pbs_errno`.

pbs_connect

connect to a PBS batch server

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>

int pbs_connect(char *server)
extern char *pbs_server;
```

DESCRIPTION

A virtual stream (TCP/IP) connection is established with the server specified by server.

This function must be called before any of the other pbs_ functions. They will transmit their batch requests over the connection established by this function. Multiple requests may be issued over the connection before it is closed.

The connection should be closed by a call to pbs_disconnect() when all requests have been sent to the server.

The parameter called server is of the form

```
host_name[:port].
```

If port is not specified, the standard PBS port number will be used.

If the parameter, server, is either the null string or a null pointer, a connection will be opened to the default server. The default server is defined by (a) the setting of the environment variable PBS_DEFAULT which contains a destination, or (b) by adding the parameter PBS_SERVER to the global configuration file /etc/pbs.conf.

The variable pbs_server, declared in pbs_ifl.h, is set on return to point to the server name to which pbs_connect() connected or attempted to connect.

pbs_connect() determines whether or not the complex has a failover server configured. It also determines which server is the primary and which is the secondary. pbs_connect() is called by client commands, and directs traffic to the correct server.

In order to use `pbs_connect` with Windows, initialize the network library and link with `winsock2`. To do this, call `winsock_init()` before calling `pbs_connect()`, and link against the `ws2_32.lib` library.

SEE ALSO

`qsub(1B)`, `pbs_alterjob(3B)`, `pbs_deljob(3B)`, `pbs_disconnect(3B)`,
`pbs_geterrmsg(3B)`, `pbs_holdjob(3B)`, `pbs_locate(3B)`,
`pbs_manager(3B)`, `pbs_movejob(3B)`, `pbs_msgjob(3B)`,
`pbs_rerunjob(3B)`, `pbs_rlsjob(3B)`, `pbs_runjob(3B)`,
`pbs_selectjob(3B)`, `pbs_selstat(3B)`, `pbs_sigjob(3B)`,
`pbs_statjob(3B)`, `pbs_statque(3B)`, `pbs_statserver(3B)`,
`pbs_submit(3B)`, `pbs_terminate(3B)`, `pbs_server(8B)`, and the PBS
Professional Programmer's Guide

DIAGNOSTICS

When the connection to batch server has been successfully created, the routine will return a connection identifier which is positive. Otherwise, a negative value is returned. The error number is set in `pbs_erno`.

pbs_default

return the name of the default PBS server

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
char *pbs_default()
```

DESCRIPTION

A character string is returned containing the name of the default PBS server. The default server is defined by (a) the setting of the environment variable `PBS_DEFAULT` which contains a destination, or (b) by adding the parameter `PBS_SERVER` to the global configuration file `/etc/pbs.conf`.

DIAGNOSTICS

If the default server cannot be determined, a NULL value is returned.

SEE ALSO

`qsub(1B)`, `pbs_connect(3B)`, `pbs_disconnect(3B)`, and the PBS Professional Programmer's Guide

pbs_deljob

delete a PBS batch job

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_deljob(int connect, char *job_id, char *extend)
```

DESCRIPTION

Issue a batch request to delete a batch job. If the batch job is running, the execution server will send the SIGTERM signal followed by SIGKILL.

A Delete Job batch request is generated and sent to the server over the connection specified by connect which is the return value of pbs_connect().

The argument, job_id, identifies which job is to be deleted. It is specified in the form:

```
“sequence_number.server”
```

The argument, extend, is overloaded to serve more than one purpose. If extend points to a string other than the above, it is taken as text to be appended to the message mailed to the job owner. This mailing occurs if the job is deleted by a user other than the job owner.

SEE ALSO

qdel(1B) and pbs_connect(3B)

DIAGNOSTICS

When the batch request generated by the pbs_deljob() function has been completed successfully by a batch server, the routine will return 0 (zero). Otherwise, a non zero error is returned. The error number is also set in pbs_errno.

pbs_delresv

delete a reservation

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_delresv(int connect, char *resv_id, char *extend)
```

DESCRIPTION

Issue a batch request to delete a reservation. If the reservation is in state RESV_RUNNING, and there are jobs remaining in the reservation queue, the jobs will be deleted before the reservation is deleted.

A Delete Reservation batch request is generated and sent to the server over the connection specified by connect which is the return value of pbs_connect().

The argument, resv_id, identifies which reservation is to be deleted, it is specified in the form:

```
“R<sequence_number>.<server>”
```

The argument, extend is currently unused.

SEE ALSO

pbs_rdel(1B) and pbs_connect(3B)

DIAGNOSTICS

When the batch request generated by the pbs_delresv() function has been completed successfully by a batch server, the routine will return 0 (zero). Otherwise, a non zero error is returned. The error number is also set in pbs_errno.

pbs_disconnect

disconnect from a pbs batch server

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_disconnect(int connect)
```

DESCRIPTION

The virtual stream connection specified by connect, which was established with a server by a call to pbs_connect(), is closed.

SEE ALSO

pbs_connect(3B)

DIAGNOSTICS

When the connection to batch server has been successfully closed, the routine will return zero. Otherwise, a non zero error is returned. The error number is also set in pbs_erno.

pbs_geterrmsg

get error message for last pbs batch operation

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
char *pbs_geterrmsg(int connect)
```

DESCRIPTION

Return the error message text associated with a batch server request.

If the preceding batch interface library call over the connection specified by `connect` resulted in an error return from the server, there may be an associated text message. If it exists, this function will return a pointer to the null terminated text string.

SEE ALSO

`pbs_connect(3B)`

DIAGNOSTICS

If an error text message was returned by a server in reply to the previous call to a batch interface library function, `pbs_geterrmsg()` will return a pointer to it. Otherwise, `pbs_geterrmsg()` returns the null pointer.

pbs_holdjob

place a hold on a pbs batch job

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_holdjob(int connect, char *job_id, char *hold_type,
char *extend)
```

DESCRIPTION

Issue a batch request to place a hold upon a job.

A Hold Job batch request is generated and sent to the server over the connection specified by connect which is the return value of pbs_connect().

The argument, job_id , identifies which job is to be held, it is specified in the form:

```
“sequence_number.server”
```

The parameter, hold_type , contains the type of hold to be applied. The possible values are defined in pbs_ifl.h.

If hold_type is either a null pointer or points to a null string, USER_HOLD will be applied.

The parameter, extend , is reserved for implementation defined extensions.

SEE ALSO

qhold(1B), pbs_connect(3B), pbs_alterjob(3B), and pbs_rlsjob(3B)

DIAGNOSTICS

When the batch request generated by pbs_holdjob () function has been completed successfully by a batch server, the routine will return 0 (zero). Otherwise, a non zero error is returned. The error number is also set in pbs_errno.

pbs_locjob

locate current location of a pbs batch job

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
char *pbs_locjob(int connect, char *job_id, char *extend)
```

DESCRIPTION

Issue a batch request to locate a batch job. If the server currently manages the batch job, or knows which server does currently manage the job, it will reply with the location of the job.

A Locate Job batch request is generated and sent to the server over the connection specified by `connect` which is the return value of `pbs_connect()`.

The argument, `job_id`, identifies which job is to be located, it is specified in the form:

```
“sequence_number.server”
```

The argument, `extend`, is reserved for implementation defined extensions. It is not currently used by this function.

The return value is a pointer to a character string which contains the current location if known. The syntax of the location string is:

```
“server_name”.
```

If the location of the job is not known, the return value is the NULL pointer.

SEE ALSO

`qsub(1B)` and `pbs_connect(3B)`

DIAGNOSTICS

When the batch request generated by the `pbs_locjob()` function has been completed successfully by a batch server, the routine will return a non null pointer to the destination. Otherwise, a null pointer is returned. The error number is set in `pbs_erno`.

pbs_manager

modifies a PBS batch object

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_manager(int connect, int command, int obj_type, char *obj_name,
struct attrpl *attrib, char *extend)
```

DESCRIPTION

Issue a batch request to perform administration functions at a server. With this request, server objects such as queues can be created and deleted, and have their attributes set and unset.

A Manage batch request is generated and sent to the server over the connection specified by `connect` which is the return value of `pbs_connect()`. This request requires full batch administrator privilege.

The `parameter, command`, specifies the operation to be performed. See `pbs_ifl.h`:

```
MGR_CMD_CREATE  creates the object
MGR_CMD_DELETE  deletes the object
MGR_CMD_SET     sets the value
MGR_CMD_UNSET   unsets the value
MGR_CMD_IMPORT  imports the hook
MGR_CMD_EXPORT  exports the hook
```

The `parameter, obj_type`, declares the type of object upon which the command operates. See `pbs_ifl.h`:

```
MGR_OBJ_SERVER  Server object
MGR_OBJ_QUEUE   Queue object
MGR_OBJ_NODE    Node object
MGR_OBJ_HOOK    Hook object
```

The `parameter, obj_name`, is the name of the specific object.

The `parameter, attrib`, is a pointer to an `attrpl` structure which is defined in `pbs_ifl.h` as:

```

struct attrpl {
    char *name;
    char *resource;
    char *value;
    enum batch_op op;
    struct attrpl *next;
};

```

The attr list is terminated by the first entry where next is a null pointer.

The name member points to a string which is the name of the attribute.

If the attribute is one which contains a set of resources, the specific resource is specified in the structure member resource. Otherwise, the member resource is pointer to a null string.

The value member points to a string which is the new value of the attribute. For parameterized limit attributes, this string contains all parameters for the attribute.

The op member defines the manner in which the new value is assigned to the attribute. The operators are:

```
“enum batch_op { ..., SET, UNSET, INCR, DECR };”
```

For MGR_CMD_IMPORT, specify attrpl “name” as “content-type”, “content-encoding”, and “input-file” along with the corresponding “value” and an “op” of SET.

For MGR_CMD_EXPORT, specify attrpl “name” as “content-type”, “content-encoding”, and “output-file” along with the corresponding “value” and an “op” of SET.

The parameter extend is reserved for implementation-defined extensions.

Privilege required for functions depends on whether those functions are used with hooks.

When not used with hooks:

Functions MGR_CMD_CREATE and MGR_CMD_DELETE require PBS Manager privilege.

Functions `MGR_CMD_SET` and `MGR_CMD_UNSET` require PBS Manager or Operator privilege.

When used with hooks:

All commands require root privilege on the server host.

Functions `MGR_CMD_IMPORT`, `MGR_CMD_EXPORT`, and `MGR_OBJ_HOOK` are used only with hooks, and therefore require root privilege on the server host.

DIAGNOSTICS

When the batch request generated by `pbs_manager()` function has been completed successfully by a batch server, the routine will return 0 (zero). Otherwise, a non zero error is returned. The error number is also set in `pbs_errno`.

SEE ALSO

The PBS Professional Programmer's Guide,
`qmgr(1B)`, `pbs_connect(3B)`

pbs_movejob

move a pbs batch job to a new destination

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_movejob(int connect, char *job_id, char *destination, char *extend)
```

DESCRIPTION

Issue a batch request to move a job to a new destination. The job is removed from the present queue and instantiated in a new queue.

A Move Job batch request is generated and sent to the server over the connection specified by connect which is the return value of pbs_connect().

The job_id parameter identifies which job is to be moved; it is specified in the form: "sequence_number.server"

The destination parameter specifies the new destination for the job. It is specified as: [queue][@server].

If destination is a null pointer or a null string, the destination will be the default queue at the current server. If destination specifies a queue but not a server, the destination will be the named queue at the current server. If destination specifies a server but not a queue, the destination will be the default queue at the named server. If destination specifies both a queue and a server, the destination is that queue at that server.

A job in the Running, Transiting, or Exiting state cannot be moved.

The parameter, extend, is reserved for implementation defined extensions.

SEE ALSO

qmove(1B), qsub(1B), and pbs_connect(3B)

DIAGNOSTICS

When the batch request generated by pbs_movejob() function has been completed successfully by a batch server, the routine will return 0 (zero). Otherwise, a non zero error is returned. The error number is also set in pbs_errno.

pbs_msgjob

record a message for a running pbs batch job

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_msgjob(int connect, char *job_id, int file, char *message,
char *extend)
```

DESCRIPTION

Issue a batch request to write a message in an output file of a batch job.

A Message Job batch request is generated and sent to the server over the connection specified by `connect` which is the return value of `pbs_connect()`.

The argument, `job_id`, identifies the job to which the message is to be sent; it is specified in the form:

```
“sequence_number.server”
```

The parameter, `file`, indicates the file or files to which the message string is to be written. See `pbs_ifl.h` for acceptable values.

The parameter, `message`, is the message string to be written.

The parameter, `extend`, is reserved for implementation defined extensions.

SEE ALSO

`qmsg(1B)` and `pbs_connect(3B)`

DIAGNOSTICS

When the batch request generated by `pbs_msgjob()` function has been completed successfully by a batch server, the routine will return 0 (zero). Otherwise, a non zero error is returned. The error number is also set in `pbs_errno`.

pbs_orderjob

reorder pbs batch jobs in a queue

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_orderjob(int connect, char *job_id1, char *job_id2,
char *extend)
```

DESCRIPTION

Issue a batch request to swap the order of two jobs with in a single queue.

An Order Job batch request is generated and sent to the server over the connection specified by connect which is the return value of pbs_connect().

The parameters job_id1 and job_id2 identify which jobs are to be swapped. They are specified in the form:
“sequence_number.server”.

The parameter, extend, is reserved for implementation defined extensions.

SEE ALSO

qorder(1B), qmove(1B), qsub(1M), and pbs_connect(3B)

DIAGNOSTICS

When the batch request generated by pbs_orderjob() function has been completed successfully by a batch server, the routine will return 0 (zero). Otherwise, a non zero error is returned. The error number is also set in pbs_errno.

pbs_rerunjob

rerun a pbs batch job

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_rerunjob(int connect, char *job_id, char *extend)
```

DESCRIPTION

Issue a batch request to rerun a batch job.

A Rerun Job batch request is generated and sent to the server over the connection specified by connect which is the return value of pbs_connect().

If the job is marked as being not rerunnable, the request will fail and an error will be returned.

The argument, job_id , identifies which job is to be rerun it is specified in the form:

```
“sequence_number.server”
```

The parameter, extend , is reserved for implementation defined extensions.

SEE ALSO

qrerun(1B), qsub(1B), and pbs_connect(3B)

DIAGNOSTICS

When the batch request generated by pbs_rerunjob() function has been completed successfully by a batch server, the routine will return 0 (zero). Otherwise, a non zero error is returned. The error number is also set in pbs_errno.

pbs_resreserve, pbs_resrelease

reserve/free batch resources

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_resreserve(int connect, char **resourcelist, int arraysize,
resource_t *resource_id)
```

```
int pbs_resrelease(int connect, resource_t resource_id)
```

DESCRIPTION

pbs_resreserve

Issue a request to the batch server to reserve specified resources. `connect` is the connection returned by `pbs_connect()`. `resourcelist` is an array of one or more strings specifying the resources to be queried. `arraysize` is the number of strings in `resourcelist`. `resource_id` is a pointer to a resource handle. The pointer cannot be null. If the present value of the resource handle is `RESOURCE_T_NULL`, this request is for a new reservation and if successful, a resource handle will be returned in `resource_id`.

If the value of `resource_id` as supplied by the caller is not `RESOURCE_T_NULL`, this is an existing (partial) reservation. Resources currently reserved for this handle will be released and the full reservation will be attempted again. If the caller wishes to release the resources allocated to a partial reservation, the caller should pass the resource handle to `pbs_resrelease()`.

At the present time the only resources which may be specified are "nodes". It should be specified as

```
nodes=specification
```

where `specification` is what a user specifies in the `-l` option argument list for `nodes`, see `qsub (1B)`.

pbs_resrelease

The `pbs_resrelease()` call releases or frees resources reserved with the resource handle of `resource_id` returned from a prior `pbs_resreserve()` call. `connect` is the connection returned by `pbs_connect()`.

Both functions require that the issuing user have operator or administrator privilege.

SEE ALSO

qsub(1B), pbs_connect(3B), pbs_disconnect(3B) and pbs_resources(7B)

DIAGNOSTICS

pbs_resreserve() and pbs_resrelease() return zero on success. Otherwise, a non zero error is returned. The error number is also set in pbs_erno.

PBSE_RMPART

is a special case indicating that some but not all of the requested resources could be reserved; a partial reservation was made. The reservation request should either be rerequested with the returned handle or the partial resources released.

PBSE_RMBADPARAM

a parameter is incorrect, such as a null for the pointer to the resource_id.

PBSE_RMNOPARAM

a parameter is missing, such as a null resource list.

pbs_rlsjob

release a hold on a pbs batch job

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_rlsjob(int connect, char *job_id, char *hold_type, char *extend)
```

DESCRIPTION

Issue a batch request to release a hold from a job.

A Release Job batch request is generated and sent to the server over the connection specified by `connect` which is the return value of `pbs_connect()`.

The argument, `job_id`, identifies the job from which the hold is to be released, it is specified in the form:

```
“sequence_number.server”
```

The parameter, `hold_type`, contains the type of hold to be released. The possible values are defined in `pbs_ifl.h`.

If `hold_type` is either a null pointer or points to a null string, `USER_HOLD` will be released.

The parameter, `extend`, is reserved for implementation defined extensions.

SEE ALSO

`qrls(1B)`, `qhold(1B)`, `qalter(1B)`, `pbs_alterjob(3B)`, `pbs_connect(3B)`, and `pbs_holdjob(3B)`

DIAGNOSTICS

When the batch request generated by `pbs_rlsjob()` function has been completed successfully by a batch server, the routine will return 0 (zero). Otherwise, a non zero error is returned. The error number is also set in `pbs_errno`.

pbs_runjob, pbs_asyrunjob

run a PBS batch job, asynchronous batch job

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_runjob(int connect, char *job_id, char *location, char *extend)
```

```
int pbs_asyrunjob(int connect, char *job_id, char *location, char *extend)
```

DESCRIPTION

Issue a batch request to run a batch job.

For `pbs_runjob()` a "Run Job" batch request is generated and sent to the server over the connection specified by `connect` which is the return value of `pbs_connect()`. The server will reply when the job has started execution unless file in-staging is required. In that case, the server will reply when the staging operations are started.

For `pbs_asyrunjob()` an "Asynchronous Run Job" request is generated and sent to the server over the connection. The server will validate the request and reply before initiating the execution of the job. This version of the call can be used to reduce latency in scheduling, especially when the scheduler must start a large number of jobs.

These requests requires that the issuing user have operator or administrator privilege.

The argument, `job_id`, identifies which job is to be run it is specified in the form:

sequence_number.server

The argument, `location`, if not the null pointer or null string, specifies the location where the job should be run, and optionally the resources to use. The location is the same as the `-H` option to the `qrun` command. See the description of `qrun -H`, both with and without resources specified, in the `qrun.8B` man page.

The argument, `extend`, is reserved for implementation-defined extensions.

SEE ALSO

qrun(8B), qsub(1B), and pbs_connect(3B)

DIAGNOSTICS

When the batch request generated by the `pbs_runjob()` or `pbs_asyruntime()` functions has been completed successfully by a batch server, the routines will return 0 (zero). Otherwise, a non zero error is returned.

The error number is also set in `pbs_erno`.

pbs_selectjob

select pbs batch jobs

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
char **pbs_selectjob(int connect, struct attrpl *attrib, char *extend)
```

DESCRIPTION

Issue a batch request to select jobs which meet certain criteria.

pbs_selectjob() returns an array of job identifiers which met the criteria.

The attrpl struct contains the list of selection criteria.

Initially all batch jobs are selected for which the user is authorized to query status. This set may be reduced or filtered by specifying certain attributes of the jobs.

A Select Jobs batch request is generated and sent to the server over the connection specified by connect which is the return value of pbs_connect().

The argument, attrib, is a pointer to an attrpl structure which is defined in pbs_ifl.h as:

```
struct attrpl {
    struct attrpl *next;
    char *name;
    char *resource;
    char *value;
    enum batch_op op;
};
```

The attrib list is terminated by the first entry where next is a null pointer.

The name member points to a string which is the name of the attribute. Not all of the job attributes may be used as a selection criteria. The resource member points to a string which is the name of a resource. This member is only used when name is set to ATTR_I. Otherwise, resource should be a pointer to a null string. The value member points to a string which is the value of the attribute or resource. The attribute names are listed in pbs_job_attributes.7B.

The op member defines the operator in the logical expression:

```
value operator current_value
```

The logical expression must evaluate as true for the job to be selected. The permissible values of `op` are defined in `pbs_ifl.h` as:

```
“enum batch_op { ..., EQ, NE, GE, GT, LE, LT, ... };”
```

The attributes marked with (E) in the description above may only be selected with the equal, EQ, or not equal, NE, operators.

If `attrib` itself is a null pointer, then no selection is done on the basis of attributes.

The return value is a pointer to a null terminated array of character pointers. Each character pointer in the array points to a character string which is a job identifier in the form:

```
sequence_number.server@server
```

The array is allocated by `pbs_selectjob` via `malloc()`. When the array is no longer needed, the user is responsible for freeing it by a call to `free()`.

The parameter, `extend`, is reserved for implementation defined extensions.

Finished and Moved Jobs

In order to get information on finished and moved jobs, you must add an ‘x’ character to the `extend` parameter. The `extend` parameter is a character string; set one character to be the ‘x’ character. For example:

```
pbs_selectjob ( ..., ..., extend) ...
```

To get information on finished and moved jobs only, specify the Finished (‘F’) and moved (‘M’) job states. You must also use the `extend` character string containing the ‘x’ character.

Subjobs are not considered finished until the parent array job is finished.

SEE ALSO

`qselect(1B)`, `pbs_alterjob(3B)`, and `pbs_connect(3B)`

DIAGNOSTICS

When the batch request generated by `pbs_selectjob()` function has been completed successfully by a batch server, the routine will return a pointer to the array of job identifiers. If no jobs met the criteria, the first pointer in the array will be the null pointer.

If an error occurred, a null pointer is returned and the error is available in the global integer `pbs_erno`.

pbs_selstat

obtain status of selected pbs batch jobs

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
struct batch_status *pbs_selstat(int connect, struct attrpl *sel_list,
struct attrl *rattrib char *extend)
```

```
void pbs_statfree(struct batch_status *psj)
```

DESCRIPTION

Issue a batch request to examine the status of jobs which meet certain criteria. `pbs_selstat()` returns a list of `batch_status` structures for those jobs which met the selection criteria.

The `sel_list` struct holds the selection criteria. The `rattrib` struct holds the list of attributes whose values are to be returned.

This function is a combination of `pbs_selectjobs()` and `pbs_statjob()`. It is an extension to the POSIX Batch standard.

Initially all batch jobs are selected for which the user is authorized to query status. This set may be reduced or filtered by specifying certain attributes of the jobs.

A Select Status batch request is generated and sent to the server over the connection specified by `connect` which is the return value of `pbs_connect()`.

The parameter, `sel_list`, is a pointer to an `attrpl` structure which is defined in `pbs_ifl.h` as:

```
struct attrpl {
    struct attrpl *next;
    char *name;
    char *resource;
    char *value;
    enum batch_op op;
};
```

The `sel_list` list is terminated by the first entry where `next` is a null pointer.

The `name` member points to a string which is the name of the attribute. Not all of the job attributes may be used as a selection criteria. The `resource` member points to a string which is the name of a resource. This member is only used when `name` is set to `ATTR_I`, otherwise it should be a pointer to a null string. The `value` member points to a string which is the value of the attribute or resource. The attribute names are listed in `pbs_job_attributes.7B`.

The `op` member defines the operator in the logical expression:

```
value operator current_value
```

The logical expression must evaluate as true for the job to be selected. The permissible values of `op` are defined in `pbs_ifl.h` as:

```
“enum batch_op { ..., EQ, NE, GE, GT, LE, LT, ... };”
```

The attributes marked with (E) in the description above may only be selected with the equal, EQ, or not equal, NE, operators.

If `sel_list` itself is a null pointer, then no selection is done on the basis of attributes.

The `parameter`, `rattrib`, is a pointer to an `attrl` structure which is defined below. The `rattrib` list is terminated by the first entry where `next` is a null pointer. If `attrib` is given, then only the attributes in the list are returned by the server. Otherwise, all the attributes of a job are returned. When an `attrib` list is specified, the `name` member is a pointer to an attribute name as listed in `pbs_alter(3)` and `pbs_submit(3)`. The `resource` member is only used if the `name` member is `ATTR_I`, otherwise it should be a pointer to a null string. The `value` member should always be a pointer to a null string.

The `return` value is a pointer to a list of `batch_status` structures or the null pointer if no jobs can be queried for status. The `batch_status` structure is defined in `pbs_ifl.h` as

```
struct batch_status {
    struct batch_status *next;
    char                *name;
    struct attrl        *attribs;
    char                *text;
}
```

The entry, `attribs`, is a pointer to a list of `attrl` structures defined in `pbs_ifl.h` as:

```

struct attrl {
    struct attrl *next;
    char      *name;
    char      *resource;
    char      *value;
};

```

It is up to the user to free the list of `batch_status` structures when no longer needed, by calling `pbs_statfree()`.

The `extend` parameter is for optional features and or additions. Normally, this should be null pointer.

When `pbs_selstat` is used to retrieve the `Submit_arguments` job attribute, PBS returns an XML-encoded string value according to the HPCBP specification.

Finished and Moved Jobs

In order to get information on finished and moved jobs, you must add an 'x' character to the `extend` parameter. The `extend` parameter is a character string; set one character to be the 'x' character. For example:

```
pbs_selstat ( ..., ..., ..., extend) ...
```

To get information on finished and moved jobs only, specify the Finished ('F') and moved ('M') job states. You must also use the `extend` character string containing the 'x' character. For example:

```

sel_list->next = sel_list;
sel_list->name = ATTR_state;
sel_list->value = "MF";
sel_list->op = EQ;
pbs_selstat ( ..., sel_list, ..., extend) ...

```

Subjobs are not considered finished until the parent array job is finished.

SEE ALSO

`qselect(1B)`, `pbs_alterjob(3B)`, `pbs_connect(3B)`, `pbs_statjob(3B)`, and `pbs_selectjob(3B)`.

DIAGNOSTICS

When the batch request generated by `pbs_selstat()` function has been completed successfully by a batch server, the routine will return a pointer to the list of `batch_status` structures. If no jobs met the criteria or an error occurred, the return will be the null pointer. If an error occurred, the global integer `pbs_erno` will be set to a non-zero value.

pbs_sigjob

send a signal to a pbs batch job

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_sigjob(int connect, char *job_id, char *signal, char *extend)
```

DESCRIPTION

Issue a batch request to send a signal to a batch job.

A Signal Job batch request is generated and sent to the server over the connection specified by `connect` which is the return value of `pbs_connect()`. If the batch job is in the running state, the batch server will send the job the signal number corresponding to the signal named in `signal`.

The argument, `job_id`, identifies which job is to be signaled, it is specified in the form:

```
“sequence_number.server”
```

The signal argument is the name of a signal. It may be the alphabetic form with or without the SIG prefix, or it may be a numeric string for the signal number. Two special names are recognized, `suspend` and `resume`. If the name of the signal is not a recognized signal name on the execution host, no signal is sent and an error is returned. If the job is not in the running state, no signal is sent and an error is returned, except when the signal is `resume` and the job is suspended.

The parameter, `extend`, is reserved for implementation defined extensions.

SEE ALSO

`qsig(1B)` and `pbs_connect(3B)`

DIAGNOSTICS

When the batch request generated by `pbs_sigjob()` function has been completed successfully by a batch server, the routine will return 0 (zero). Otherwise, a non zero error is returned. The error number is also set in `pbs_erro`.

pbs_stagein

request that files for a pbs batch job be staged in.

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_stagein(int connect, char *job_id, char *location, char *extend)
```

DESCRIPTION

Issue a batch request to start the stage in of files specified in the stagein attribute of a batch job.

A stage in batch request is generated and sent to the server over the connection specified by connect which is the return value of pbs_connect().

This request directs the server to begin the stage in of files specified in the job's stage in attribute. This request requires that the issuing user have operator or administrator privilege.

The argument, job_id , identifies which job for which file staging is to begin. It is specified in the form:

```
“sequence_number.server”
```

The argument, location , if not the null pointer or null string, specifies the location where the job will be run and hence to where the files will be staged. The location is the name of a host in the cluster managed by the server. If the job is then directed to run at different location, the run request will be rejected.

The argument, extend , is reserved for implementation defined extensions.

SEE ALSO

qrun(8B), qsub(1B), and pbs_connect(3B)

DIAGNOSTICS

When the batch request generated by pbs_stagein() function has been completed successfully by a batch server, the routine will return 0 (zero). Otherwise, a non zero error is returned. The error number is also set in pbs_errno.

pbs_statjob

obtain status of pbs batch jobs

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
struct batch_status *pbs_statjob(int connect, char *id, struct attrl *attrib, char *extend)
```

```
void pbs_statfree(struct batch_status *psj)
```

DESCRIPTION

Issue a batch request to obtain the status of a specified batch job or a set of jobs at a destination.

A Status Job batch request is generated and sent to the server over the connection specified by connect which is the return value of pbs_connect().

The parameter, id, may be either a job identifier or a destination identifier.

If id is a job identifier, it is the identifier of the job for which status is requested. It is specified in the form:

```
“sequence_number.server”
```

If id is a destination identifier, it specifies that status of all jobs at the destination (queue) which the user is authorized to see be returned. If id is the null pointer or a null string, the status of each job at the server which the user is authorized to see is returned.

The parameter, attrib, is a pointer to an attrl structure which is defined in pbs_ifl.h as:

```
struct attrl {
    struct attrl *next;
    char *name;
    char *resource;
    char *value;
};
```

The attrib list is terminated by the first entry where next is a null pointer. If attrib is given, then only the attributes in the list are returned by the server. Otherwise, all the attributes of a job are returned. When an attrib list is specified, the name member is a pointer to an attribute name as listed in pbs_alter(3) and pbs_submit(3). The resource member is only used if the name member is ATTR_1, otherwise it should be a pointer to a null string. The

value member should always be a pointer to a null string.

The parameter, `extend`, is reserved for implementation defined extensions.

The return value is a pointer to a list of `batch_status` structures or the null pointer if no jobs can be queried for status. The `batch_status` structure is defined in `pbs_ifl.h` as

```
struct batch_status {
    struct batch_status *next;
    char                *name;
    struct attrl        *attribs;
    char                *text;
}
```

It is up to the user to free the structure when no longer needed, by calling `pbs_statfree()`.

When `pbs_statjob` is used to retrieve the `Submit_arguments` job attribute, PBS returns an XML-encoded string value according to the HPCBP specification.

Finished and Moved Jobs

When querying for multiple jobs, to get information on finished and moved jobs, you must add an 'x' character to the `extend` parameter. The `extend` parameter is a character string; set one character to be the 'x' character.

When querying for multiple jobs, to get information on finished and moved jobs only, specify the Finished ('F') and moved ('M') job states. You must also use the `extend` character string containing the 'x' character.

When querying for a single finished job, the `extend` string does not need to contain the 'x' character.

Subjobs are not considered finished until the parent array job is finished.

SEE ALSO

`qstat(1B)` and `pbs_connect(3B)`

DIAGNOSTICS

When the batch request generated by `pbs_statjob()` function has been completed successfully and the status of each job has been returned by the batch server, the routine will return a pointer to the list of `batch_status` structures. If no jobs were available to query or an error occurred, a null pointer is returned. The global integer `pbs_errno` should be examined to determine the cause.

pbs_statnode, pbs_statvnode, pbs_stathost

obtain status of PBS vnodes or hosts

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>

struct batch_status *pbs_stathost(int connect, char *id,
struct attrl *attrib, char *extend)

struct batch_status *pbs_statnode(int connect, char *id,
struct attrl *attrib, char *extend)

struct batch_status *pbs_statvnode(int connect, char *id,
struct attrl *attrib, char *extend)

void pbs_statfree(struct batch_status *psj)
```

DESCRIPTION

Issue a batch request to obtain the status of PBS execution hosts or vnodes.

`pbs_stathost` returns information about the single host named in the call or about all hosts known to the PBS Server.

`pbs_statnode` is identical to `pbs_stathost` in function. It is retained for backward compatibility.

`pbs_statvnode` returns information about the single virtual node (vnode) named in the call or about all vnodes known to the PBS Server.

A Status Node batch request is generated and sent to the server over the connection specified by `connect` which is the return value of `pbs_connect()`.

The `id` is the name of a host for `pbs_stathost`, or a vnode for `pbs_statvnode`, or the null string. If `id` specifies a name, the status of that host or vnode will be returned. If the `id` is a null string (or null pointer), the status of all hosts or vnodes at the server will be returned.

The parameter, `attrib`, is a pointer to an `attrl` structure which is defined in `pbs_ifl.h` as:

```

struct attrl {
    struct attrl *next;
    char    *name;
    char    *resource;
    char    *value;
};

```

The attrib list is terminated by the first entry where next is a null pointer. If attrib is given, then only the attributes in the list are returned by the server. Otherwise, all the attributes of a node are returned. When an attrib list is specified, the name member is a pointer to an attribute name. The resource member is not used and must be a pointer to a null string. The value member should always be a pointer to a null string.

The parameter, extend, is reserved for implementation defined extensions.

The return value is a pointer to a list of batch_status structures, which is defined in pbs_ifl.h as:

```

struct batch_status {
    struct batch_status *next;
    char    *name;
    struct attrl    *attribs;
    char    *text;
}

```

It is up to the user to free the structure when no longer needed, by calling pbs_statfree().

DIAGNOSTICS

When the batch request generated by pbs_stathost(), pbs_statnode(), or pbs_statvnode() function has been completed successfully by a batch server, the routine will return a pointer to the batch_status structure. Otherwise, a null pointer is returned and the error code is set in the global integer pbs_erno.

SEE ALSO

qstat(1B), pbs_connect(3B)

pbs_statque

obtain status of pbs batch queues

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>

struct batch_status *pbs_statque(int connect, char *id,
struct attrl *attrib, char *extend)

void pbs_statfree(struct batch_status *psj)
```

DESCRIPTION

Issue a batch request to obtain the status of a batch queue.

A Status Queue batch request is generated and sent to the server over the connection specified by connect which is the return value of pbs_connect().

The id is the name of a queue, in the form:

queue_name

or the null string. If

queue_name

is specified, the status of the queue named

queue_name

at the server is returned. If the id is a null string or null pointer, the status of all queues at the server is returned.

The parameter, attrib, is a pointer to an attrl structure which is defined in pbs_ifl.h as:

```
struct attrl {
    struct attrl *next;
    char *name;
    char *resource;
    char *value;
};
```

The attrib list is terminated by the first entry where next is a null pointer. If attrib is given, then only the attributes in the list are returned by the server. Otherwise, all the attributes of a queue are

returned. When an attrib list is specified, the name member is a pointer to an attribute name as listed in `pbs_alterjob(3B)` and `pbs_submit(3B)`. The resource member is only used if the name member is `ATTR_1`, otherwise it should be a pointer to a null string. The value member should always be a pointer to a null string.

When `pbs_statque` is used to get the attributes of an object, a single attrl data structure is returned for each parameterized attribute.

The parameter, `extend`, is reserved for implementation defined extensions.

The return value is a pointer to a list of `batch_status` structures, which is defined in `pbs_ifl.h` as:

```
struct batch_status {
    struct batch_status *next;
    char                *name;
    struct attrl        *attribs;
    char                *text;
}
```

It is up the user to free the structure when no longer needed, by calling `pbs_statfree()`.

SEE ALSO

`qstat(1B)` and `pbs_connect(3B)`

DIAGNOSTICS

When the batch request generated by `pbs_statque()` function has been completed successfully by a batch server, the routine will return a pointer to the `batch_status` structure. Otherwise, a null pointer is returned and the error code is set in the global integer `pbs_erno`.

pbs_statresv

obtain status information about reservations

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>

struct batch_status *pbs_statresv(int connect, char *id,
struct attrl *attrib, char *extend)

void pbs_statfree(struct batch_status *psj)
```

DESCRIPTION

Issue a batch request to obtain the status of a specified reservation or a set of reservations at a destination.

A Status Reservation batch request is generated and sent to the server over the connection specified by `connect` which is the return value of `pbs_connect()`.

The parameter, `id`, is a reservation identifier. A reservation identifier is of the form:

```
“R<sequence_number>.<server>”
```

If `id` is the null pointer or a null string, the status of each reservation at the server which the user is authorized to see is returned.

The parameter, `attrib`, is a pointer to an `attrl` structure which is defined in `pbs_ifl.h` as:

```
struct attrl {
    struct attrl *next;
    char *name;
    char *resource;
    char *value;
};
```

The `attrib` list is terminated by the first entry where `next` is a null pointer. If `attrib` is given, then only the attributes in the list are returned by the server. Otherwise, all the attributes of a reservation are returned. When an `attrib` list is specified, the `name` member is a

pointer to a attribute name as listed in `pbs_submitresv(3)`. The resource member is only used if the name member is `ATTR_1`, otherwise it should be a pointer to a null string. The value member should always be a pointer to a null string.

The parameter, `extend`, is reserved for implementation defined extensions.

The return value is a pointer to a list of `batch_status` structures or the null pointer if no reservations can be queried for status. The `batch_status` structure is defined in `pbs_ifl.h` as

```
struct batch_status {
    struct batch_status *next;
    char                *name;
    struct attrl        *attribs;
    char                *text;
}
```

It is up the user to free the structure when no longer needed, by calling `pbs_statfree()`.

SEE ALSO

`pbs_rstat(1B)` and `pbs_connect(3B)`

DIAGNOSTICS

When the batch request generated by `pbs_statresv()` function has been completed successfully and the status of each reservation has been returned by the batch server, the routine will return a pointer to the list of `batch_status` structures. If no reservations were available to query or an error occurred, a null pointer is returned. The global integer `pbs_errno` should be examined to determine the cause.

pbs_statsched

obtain status of PBS scheduler

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
struct batch_status *pbs_statsched(int connect, struct attrl *attrib,
char *extend)
```

```
void pbs_statfree(struct batch_status *psj)
```

DESCRIPTION

Issue a batch request to obtain the status of PBS scheduler.

A Status Scheduler batch request is generated and sent to the server. The parameter `connect` is the return value of `pbs_connect()`.

The parameter, `attrib`, is a pointer to an `attrl` structure which is defined in `pbs_ifl.h` as:

```
struct attrl {
    struct attrl *next;
    char *name;
    char *resource;
    char *value;
};
```

The `attrib` list is terminated by the first entry where `next` is a null pointer. If `attrib` is given, then only the attributes in the list are returned by the server. Otherwise, all the attributes of the scheduler are returned. When an `attrib` list is specified, the `name` member is a pointer to an attribute name as listed in `pbs_alter(3)` and `pbs_submit(3)`. The `resource` member is only used if the `name` member is `ATTR_I`, otherwise it should be a pointer to a null string. The `value` member should always be a pointer to a null string.

The parameter, `extend`, is reserved for implementation-defined extensions.

The return value of `pbs_statsched()` is a pointer to a list of

batch_status structures, which is defined in pbs_ifl.h as:

```
struct batch_status {
    struct batch_status *next;
    char                *name;
    struct attrl        *attribs;
    char                *text;
}
```

It is up the user to free the batch_status structure when it is no longer needed, by calling pbs_statfree().

SEE ALSO

qstat(1B) and pbs_connect(3B)

DIAGNOSTICS

When the batch request generated by pbs_statsched() has been completed successfully by the PBS server, pbs_statsched() will return a pointer to a batch_status structure. Otherwise, a null pointer is returned and the error code is set in pbs_errno.

pbs_statsserver

obtain status of a pbs batch server

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
struct batch_status *pbs_statsserver(int connect, struct attrl *attrib,
char *extend)
```

```
void pbs_statfree(struct batch_status *psj)
```

DESCRIPTION

Issue a batch request to obtain the status of a batch server.

A Status Server batch request is generated and sent to the server over the connection specified by `connect` which is the return value of `pbs_connect()`.

The parameter, `attrib`, is a pointer to an `attrl` structure which is defined in `pbs_ifl.h` as:

```
struct attrl {
    struct attrl *next;
    char *name;
    char *resource;
    char *value;
};
```

The `attrib` list is terminated by the first entry where `next` is a null pointer. If `attrib` is given, then only the attributes in the list are returned by the server. Otherwise, all the attributes of the server are returned. When an `attrib` list is specified, the `name` member is a pointer to an attribute name as listed in `pbs_alterjob(3B)` and `pbs_submit(3B)`. The `resource` member is only used if the `name` member is `ATTR_1`, otherwise it should be a pointer to a null string. The `value` member should always be a pointer to a null string.

When `pbs_statsserver` is used to get the attributes of an object, a single `attrl` data structure is returned for each parameterized attribute.

The parameter, `extend`, is reserved for implementation defined extensions.

The return value is a pointer to a list of `batch_status` structures, which is defined in `pbs_ifl.h` as:

```
struct batch_status {
    struct batch_status *next;
    char                *name;
    struct attrl        *attrs;
    char                *text;
}
```

It is up to the user to free the space when no longer needed, by calling `pbs_statfree()`.

SEE ALSO

`qstat(1B)` and `pbs_connect(3B)`

DIAGNOSTICS

When the batch request generated by `pbs_statserver()` function has been completed successfully by a batch server, the routine will return a pointer to a `batch_status` structure. Otherwise, a null pointer is returned and the error code is set in `pbs_errno`.

pbs_submit

submit a pbs batch job

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>

char *pbs_submit(int connect, struct attrpl *attrib,
char *script, char *destination, char *extend)
```

DESCRIPTION

Issue a batch request to submit a new batch job.

A Queue Job batch request is generated and sent to the server over the connection specified by `connect` which is the return value of `pbs_connect()`. The job will be submitted to the queue specified by `destination`.

The parameter, `attrib`, is a list of `attrpl` structures which is defined in `pbs_ifl.h` as:

```
struct attrl {
    char *name;
    char *resource;
    char *value;
    struct attrl *next;
    enum batch_op op;
};
```

The `attrib` list is terminated by the first entry where `next` is a null pointer.

The `name` member points to a string which is the name of the attribute. The `value` member points to a string which is the value of the attribute. The attribute names are defined in `pbs_job_attributes(7B)`.

If an attribute is not named in the `attrib` array, the default action will be taken. It will either be assigned the default value or will not be passed with the job. The action depends on the attribute. If `attrib` itself is a null pointer, then the default action will be taken for each attribute.

Associated with an attribute of type ATTR_1 (the letter ell) is a resource name indicated by resource in the attrl structure. All other attribute types should have a pointer to a null string for resource .

The op member is forced to a value of SET by pbs_submit().

The parameter, script , is the path name to the job script. If the path name is relative, it will be expanded to the processes current working directory. If script is a null pointer or the path name pointed to is specified as the null string, no script is passed with the job.

The destination parameter specifies the destination for the job. It is specified as:

[queue]

If destination is the null string or the queue is not specified, the destination will be the default queue at the connected server.

The parameter, extend , is reserved for implementation defined extensions.

The return value is a character string which is the job_identifier assigned to the job by the server. The space for the job_identifier string is allocated by pbs_submit() and should be released via a call to free() by the user when no longer needed.

SEE ALSO

qsub(1B) and pbs_connect(3B)

DIAGNOSTICS

When the batch request generated by pbs_submit() function has been completed successfully by a batch server, the routine will return a pointer to a character string which is the job identifier of the submitted batch job. Otherwise, a null pointer is returned and the error code is set in pbs_error.

pbs_submitresv

submit a pbs reservation

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
char *pbs_submitresv(int connect, struct attrpl *attrib, char *extend)
```

DESCRIPTION

Issue a batch request to submit a new reservation.

A Submit Reservation batch request is generated and sent to the server over the connection specified by `connect` which is the return value of `pbs_connect()`.

The parameter, `attrib`, is a list of `attrpl` structures which is defined in `pbs_ifl.h` as:

```
struct attrl {
    char *name;
    char *resource;
    char *value;
    struct attrl *next;
    enum batch_op op;
};
```

The `attrib` list is terminated by the first entry where `next` is a null pointer.

The `name` member points to a string which is the name of the attribute. The `value` member points to a string which is the value of the attribute. The attribute names are defined in `pbs_ifl.h`.

If an attribute is not named in the `attrib` array, the default action will be taken. It will either be assigned the default value or will not be passed with the reservation. The action depends on the attribute. If `attrib` itself is a null pointer, then the default action will be taken for each attribute.

Associated with an attribute of type `ATTR_1` (the letter ell) is a

resource name indicated by resource in the attrl structure. All other attribute types should have a pointer to a null string for resource .

The op member is forced to a value of SET by pbs_submitresv().

The parameter, extend , is reserved for implementation defined extensions.

The return value is a character string which is the reservation_identifier assigned to the job by the server. The space for the reservation_identifier string is allocated by pbs_submitresv() and should be released via a call to free() by the user when no longer needed.

SEE ALSO

pbs_rsub(1B) and pbs_connect(3B)

DIAGNOSTICS

When the batch request generated by pbs_submitresv() function has been completed successfully by a batch server, the routine will return a pointer to a character string which is the job identifier of the submitted batch job. Otherwise, a null pointer is returned and the error code is set in pbs_error.

pbs_terminate

terminate a pbs batch server

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>
```

```
int pbs_terminate(int connect, int manner, char *extend)
```

DESCRIPTION

Issue a batch request to shut down a batch server. This request requires the privilege level usually reserved for batch operators and administrators.

A Server Shutdown batch request is generated and sent to the server over the connection specified by `connect` which is the return value of `pbs_connect()`.

The `manner` parameter, `manner`, specifies the manner in which the server is shut down. The available manners are defined in `pbs_ifl.h`.

The server will not respond to the batch request until the server has completed its termination procedure.

The `extend` parameter, `extend`, is reserved for implementation defined extensions.

This call requires PBS Operator or Manager privilege.

SEE ALSO

`qterm(8B)` and `pbs_connect(3B)`

DIAGNOSTICS

When the batch request generated by `pbs_terminate()` function has been completed successfully by a batch server, the routine will return 0 (zero). Otherwise, a non zero error is returned. The error number is also set in `pbs_errno`.

Chapter 5

RPP Library

This chapter discusses the Reliable Packet Protocol (RPP) used by PBS. These functions provide reliable, flow-controlled, two-way transmission of data. Each data path will be called a "stream" in this document. The advantage of RPP over TCP is that many streams can be multiplexed over one socket. This allows simultaneous connections over many streams without regard to the system imposed file descriptor limit.

5.1 RPP Library Routines

The following manual pages document the application programming interface provided by the RPP library.

rpp_open, rpp_bind, rpp_poll, rpp_io, rpp_read, rpp_write, rpp_close, rpp_getaddr, rpp_flush, rpp_terminate, rpp_shutdown, rpp_rcommit, rpp_wcommit, rpp_eom, rpp_getc, rpp_putc

reliable packet protocol

SYNOPSIS

```
#include <sys/types.h>
#include <netinet/in.h>
#include <rpp.h>
```

```
int rpp_open(addr)
    struct sockaddr_in *addr;
```

```
int rpp_bind(port)
    int port;
```

```
int rpp_poll()
```

```
int rpp_io()
```

```
int rpp_read(stream, buf, len)
    u_int stream;
    char *buf;
    int len;
```

```
int rpp_write(stream, buf, len)
    u_int stream;
    char *buf;
    int len;
```

```
int rpp_close(stream)
    u_int stream;
```

```
struct sockaddr_in *rpp_getaddr(stream)
    u_int stream;
```

```
int rpp_flush(stream)
    u_int stream;
```

```
int rpp_terminate()
```

```
int rpp_shutdown()
```

```
int rpp_rcommit(stream, flag)
    u_int stream;
    int flag;
```

```
int rpp_wcommit(stream, flag)
    u_int stream;
    int flag;
```

```
int rpp_eom(stream)
    u_int stream;
```

```
int rpp_getc(stream)
    u_int stream;
```

```
int rpp_putc(stream, c)
    u_int stream;
    int c;
```

DESCRIPTION

These functions provide reliable, flow-controlled, two-way transmission of data. Each data path will be called a “stream” in this document. The advantage of RPP over TCP is that many streams can be multiplexed over one socket. This allows simultaneous connections over many streams without regard to the system imposed file descriptor limit.

Data is sent and received in “messages”. A message may be of any length and is either received completely or not at all. Long messages will cause the library to use large amounts of memory in the heap by calling `malloc(3V)`.

In order to use any of the above with Windows, initialize the network library and link with `winsoc2`. To do this, call `winsoc_init()` before calling the function and link against the `ws2_32.lib` library.

`rpp_open()` initializes a new stream connection to `addr` and returns the stream identifier. This is an integer with a value greater than or equal to zero. A negative number indicates an error. In this case, `errno` will be set.

`rpp_bind()` is an initialization call which is used to bind the UDP socket used by RPP to a particular port. The file descriptor of the UDP socket used by the library is returned.

`rpp_poll()` returns the stream identifier of a stream with data to read. If no stream is ready to read, a -2 is returned. A -1 is returned if an error occurs.

`rpp_io()` processes any packets which are waiting to be sent or received over the UDP socket. This routine should be called if a section of code could be executing for more than a few (~10) seconds without calling any other rpp function. A -1 is returned if an error occurs, 0 otherwise.

`rpp_read()` transfers up to `len` characters of a message from stream into `buf`. If all of a message has been read, the return value will be less than `len`. The return value could be zero if all of a message had previously been read. A -1 is returned on error. A -2 is returned if the peer has closed its connection. If `rpp_poll()` is used to determine the stream is ready for reading, the call to `rpp_read()` will return immediately. Otherwise, the call will block waiting for a message to arrive.

`rpp_write()` adds information to the current message on a stream. The data in `buf` numbering `len` characters is transferred to the stream. The number of characters added to the stream are returned or a -1 on error. In this case, `errno` will be set. A -2 is returned if the peer has closed its connection.

`rpp_close()` disconnects the stream from its peer and frees all resources associated with the stream. The return value is -1 on error and 0 otherwise.

`rpp_getaddr()` returns the address which a stream is connected to. If the stream is not open, a NULL pointer is returned.

`rpp_flush()` marks the end of a message and commits all the data which has been written to the specified stream. A zero is returned if the message has been successfully committed. A -1 is returned on error.

`rpp_terminate()` is used to free all memory associated with all streams and close the UDP socket. This is done without attempting to send any final messages that may be waiting. If a process is using rpp and

calls `fork()`, the child must call `rpp_terminate()` so it will not cause a conflict with the parent's communication.

`rpp_shutdown()` is used to free all memory associated with all streams and close the UDP socket. An attempt is made to send all outstanding messages before returning.

`rpp_rcommit()` is used to "commit" or "de-commit" the information read from a message. As calls are made to `rpp_read()`, the number of characters transferred out of the message are counted. If `rpp_rcommit()` is called with `flag` being non-zero (TRUE), the current position in the message is marked as the commit point. If `rpp_rcommit()` is called with `flag` being zero (FALSE), a subsequent call to `rpp_read()` will return characters from the message following the last commit point. If an entire message has been read, `rpp_read()` will continue to return zero as the number of bytes transferred until `rpp_eom()` is called to commit the complete message.

`rpp_wcommit()` is used to "commit" or "de-commit" the information written to a stream. As calls are made to `rpp_write()`, the number of characters transferred into the message are counted. If `rpp_wcommit()` is called with `flag` being non-zero (TRUE), the current position in the message is marked as the commit point. If `rpp_wcommit()` is called with `flag` being zero (FALSE), a subsequent call to `rpp_write()` will transfer characters into the stream following the last commit point. A call to `rpp_flush()` does an automatic write commit to the current position.

`rpp_eom()` is called to terminate processing of the current message.

SEE ALSO

`tcp(4P)`, `udp(4P)`

Chapter 6

TM Library

This chapter describes the PBS Task Management library. The TM library is a set of routines used to manage multi-process, parallel, and distributed applications. The current version is an implementation of the proposed (draft) PSCHED standard sponsored by NASA. Altair has since submitted this draft to the DRAMA working group of the international Global Grid Forum standards body.

6.1 TM Library Routines

The following manual pages document the application programming interface provided by the TM library.

**tm_init, tm_nodeinfo, tm_poll, tm_notify, tm_spawn, tm_kill, tm_obit,
tm_taskinfo, tm_atnode, tm_rescinfo, tm_publish, tm_subscribe,
tm_finalize, tm_attach**

task management API

SYNOPSIS

```
#include <tm.h>

int tm_init(info, roots)
    void *info;
    struct tm_roots *roots;

int tm_nodeinfo(list, nnodes)
    tm_node_id **list;
    int *nnodes;

int tm_poll(poll_event, result_event, wait, tm_errno)
    tm_event_t poll_event;
    tm_event_t *result_event;
    int wait;
    int *tm_errno;

int tm_notify(tm_signal)
    int tm_signal;

int tm_spawn(argc, argv, envp, where, tid, event)
    int argc;
    char **argv;
    char **envp;
    tm_node_id where;
    tm_task_id *tid;
    tm_event_t *event;

int tm_kill(tid, sig, event)
    tm_task_id tid;
    int sig;
    tm_event_t *event;

int tm_obit(tid, obitval, event)
    tm_task_id tid;
```

```
int *obitval;
tm_event_t *event;

int tm_taskinfo(node, tid_list, list_size, ntasks, event)
    tm_node_id node;
    tm_task_id *tid_list;
    int list_size;
    int *ntasks;
    tm_event_t *event;

int tm_atnode(tid, node)
    tm_task_id tid;
    tm_node_id *node;

int tm_rescinfo(node, resource, len, event)
    tm_node_id node;
    char *resource;
    int len;
    tm_event_t *event;

int tm_publish(name, info, len, event)
    char *name;
    void *info;
    int len;
    tm_event_t *event;

int tm_subscribe(tid, name, info, len, info_len, event)
    tm_task_id tid;
    char *name;
    void *info;
    int len;
    int *info_len;
    tm_event_t *event;

int tm_attach(jobid, cookie, pid, tid, host, port)
    char *jobid;
    char *cookie;
    pid_t pid;
    tm_task_id *tid;
    char *host;
    int port;
```

```
int tm_finalize()
```

DESCRIPTION

These functions provide a partial implementation of the task management interface part of the PSCHED API. In PBS, MOM provides the task manager functions. This library opens a tcp socket to the MOM running on the local host and sends and receives messages using the DIS protocol (described in the PBS IDS). The tm interface can only be used by a process within a PBS job.

The PSCHED Task Management API description used to create this library was committed to paper on November 15, 1996 and was given the version number 0.1. Changes may have taken place since that time which are not reflected in this library.

The API description uses several data types that it purposefully does not define. This was done so an implementation would not be confined in the way it was written. For this specific work, the definitions follow:

```
typedef int      tm_node_id; /* job-relative node id */
#define TM_ERROR_NODE ((tm_node_id)-1)
typedef int      tm_event_t; /* > 0 for real events */
#define TM_NULL_EVENT ((tm_event_t)0)
#define TM_ERROR_EVENT ((tm_event_t)-1)
typedef unsigned long tm_task_id;
#define TM_NULL_TASK (tm_task_id)0
```

There are a number of error values defined as well: TM_SUCCESS, TM_ESYSTEM, TM_ENOEVENT, TM_ENOTCONNECTED, TM_EUNKNOWNCMD, TM_ENOTIMPLEMENTED, TM_EBADENVIRONMENT, TM_ENOTFOUND.

The functions listed here are not supported on Windows.

tm_init() initializes the library by opening a socket to the MOM on the local host and sending a TM_INIT message, then waiting for the reply. The info parameter has no use and is included to conform with the PSCHED document. The roots pointer will contain valid data after the function returns and has the following structure:

```
struct tm_roots {
    tm_task_id  tm_me;
```

```

tm_task_id  tm_parent;
int         tm_nnodes;
int         tm_ntasks;
int         tm_taskpoolid;
tm_task_id  *tm_tasklist;
};

```

- tm_me The task id of this calling task.

- tm_parent The task id of the task which spawned this task or
TM_NULL_TASK if the calling task is the initial
task started by PBS.

- tm_nnodes The number of nodes allocated to the job.

- tm_ntasks This will always be 0 for PBS.

- tm_taskpoolid PBS does not support task pools so this will always
be -1.

- tm_tasklist This will be NULL for PBS.

The tm_ntasks, tm_taskpoolid and tm_tasklist fields are not filled with data specified by the PSCHED document. PBS does not support task pools and, at this time, does not return information about current running tasks from tm_init. There is a separate call to get information for current running tasks called tm_taskinfo which is described below. The return value from tm_init is TM_SUCCESS if the library initialization was successful, or an error is returned otherwise.

tm_nodeinfo() places a pointer to a malloc'ed array of tm_node_id's in the pointer pointed at by list. The order of the tm_node_id's in list is the same as that specified to MOM in the "exec_host" attribute. The int pointed to by nnodes contains the number of nodes allocated to the job. This is information that is returned during initialization and does not require communication with MOM. If tm_init has not been called, TM_ESYSTEM is returned, otherwise TM_SUCCESS is returned.

tm_poll() is the function which will retrieve information about the task management system to locations specified when other routines request an action take place. The bookkeeping for this is done by generating an event for each action. When the task manager (MOM) sends a message that an action is complete, the event is reported by tm_poll and information is placed where the caller requested it. The argument poll_event is meant to be used to request a specific event. This implementation does not use it and it must be set to

TM_NULL_EVENT or an error is returned. Upon return, the argument `result_event` will contain a valid event number or TM_ERROR_EVENT on error. If wait is zero and there are no events to report, `result_event` is set to TM_NULL_EVENT. If wait is non-zero and there are no events to report, the function will block waiting for an event. If no local error takes place, TM_SUCCESS is returned. If an error is reported by MOM for an event, then the argument `tm_erno` will be set to an error code.

`tm_notify()` is described in the PSCHED documentation, but is not implemented for PBS yet. It will return TM_ENOTIMPLEMENTED.

`tm_spawn()` sends a message to MOM to start a new task. The node id of the host to run the task is given by `where`. The parameters `argc`, `argv` and `envp` specify the program to run and its arguments and environment very much like `exec()`. The full path of the program executable must be given by `argv[0]` and the number of elements in the `argv` array is given by `argc`. The array `envp` is NULL terminated. The argument `event` points to a `tm_event_t` variable which is filled in with an event number. When this event is returned by `tm_poll`, the `tm_task_id` pointed to by `tid` will contain the task id of the newly created task.

`tm_kill()` sends a signal specified by `sig` to the task `tid` and puts an event number in the `tm_event_t` pointed to by `event`.

`tm_obit()` creates an event which will be reported when the task `tid` exits. The `int` pointed to by `obitval` will contain the exit value of the task when the event is reported.

`tm_taskinfo()` returns the list of tasks running on the node specified by `node`. The PSCHED documentation mentions a special ability to retrieve all tasks running in the job. This is not supported by PBS. The argument `tid_list` points to an array of `tm_task_id`'s which contains `list_size` elements. Upon return, `event` will contain an event number. When this event is polled, the `int` pointed to by `ntasks` will contain the number of tasks running on the node and the array will be filled in with `tm_task_id`'s. If `ntasks` is greater than `list_size`, only `list_size` tasks will be returned.

`tm_atnode()` will place the node id where the task `tid` exists in the `tm_node_id` pointed to by `node`.

`tm_resconf()` makes a request for a string specifying the resources available on a node given by the argument `node`. The string is returned

in the buffer pointed to by resource and is terminated by a NUL character unless the number of characters of information is greater than specified by len. The resource string PBS returns is formatted as follows:

A space separated set of strings from the uname system call. The order of the strings is sysname, nodename, release, version, machine.

A comma separated set of strings giving the components of the “Resource_List” attribute of the job, preceded by a colon (:). Each component has the resource name, an equal sign, and the limit value.

tm_publish() causes len bytes of information pointed at by info to be sent to the local MOM to be saved under the name given by name.

tm_subscribe() returns a copy of the information named by name for the task given by tid. The argument info points to a buffer of size len where the information will be returned. The argument info_len will be set with the size of the published data. If this is larger than the supplied buffer, the data will have been truncated.

tm_attach() commands MOM to create a new PBS “attached task” out of a session running on MOM’s host. The jobid parameter specifies the job which is to have a new task attached. If it is NULL, the system will try to determine the correct jobid. The cookie parameter must be NULL. The pid parameter must be a non-zero process id for the process which is to be added to the job specified by jobid. If tid is non-NULL, it will be used to store the task id of the new task. The host and port parameters specify where to contact MOM. host should be NULL. The return value will be 0 if a new task has been successfully created and non-zero on error. The return value will be one of the TM error numbers defined in tm.h as follows:

TM_ESYSTEM	MOM cannot be contacted
TM_ENOTFOUND	No matching job was found
TM_ENOTIMPLEMENTED	The call is not implemented/supported
TM_ESESSION	The session specified is already attached
TM_EUSER	The calling user is not permitted to attach
TM_EOWNER	The process owner does not match the job
TM_ENOPROC	The process does not exist

tm_finalize() may be called to free any memory in use by the library and close the connection to MOM.

Chapter 7

RM Library

This chapter describes the PBS Resource Monitor library. The RM library contains functions to facilitate communication with the PBS Professional resource monitor. It is set up to make it easy to connect to several resource monitors and handle the network communication efficiently.

7.1 RM Library Routines

The following “manual” pages document the application programming interface provided by the RM library.

openrm, closerm, downrm, configrm, addreq, allreq, getreq, flushreq, activereq, fullresp

resource monitor API

SYNOPSIS

```
#include <sys/types.h>
#include <netinet/in.h>
#include <rm.h>
```

```
int openrm (host, port)
char *host;
unsigned int port;
```

```
int closerm (stream)
int stream;
```

```
int downrm (stream)
int stream;
```

```
int configrm (stream, file)
int stream;
char *file;
```

```
int addreq (stream, line)
int stream;
char *line;
```

```
int allreq (line)
char *line;
```

```
char *getreq(stream)
int stream;
```

```
int flushreq()
```

```
int activereq()
```

```
void fullresp(flag)
int flag;
```

DESCRIPTION

The resource monitor library contains functions to facilitate communication with the PBS Professional resource monitor. It is set up to make it easy to connect to several resource monitors and handle the network communication efficiently.

In all these routines, the variable `pbs_erno` will be set when an error is indicated. The lower levels of network protocol are handled by the “Data Is Strings” DIS library and the “Reliable Packet Protocol” RPP library.

`configm()` causes the resource monitor to read the file named. Deprecated.

`addreq()` begins a new message to the resource monitor if necessary. Then adds a line to the body of an outstanding command to the resource monitor.

`allreq()` begins, for each stream, a new message to the resource monitor if necessary. Then adds a line to the body of an outstanding command to the resource monitor.

`getreq()` finishes and sends any outstanding message to the resource monitor. If `fullresp()` has been called to turn off “full response” mode, the routine searches down the line to find the equal sign just before the response value. The returned string (if it is not NULL) has been allocated by `malloc` and thus `free` must be called when it is no longer needed to prevent memory leaks.

`flushreq()` finishes and sends any outstanding messages to all resource monitors. For each active resource monitor structure, it checks if any outstanding data is waiting to be sent. If there is, it is sent and the internal structure is marked to show “waiting for response”.

`fullresp()` turns on, if `flag` is true, “full response” mode where `getreq()` returns a pointer to the beginning of a line of response. This is the default. If `flag` is false, the line returned by `getreq()` is just the answer following the equal sign.

`activereq()` Returns the stream number of the next stream with something to read or a negative number (the return from `rpp_poll`) if there is no stream to read.

In order to use any of the above with Windows, initialize the network library and link with `winsoc2`. To do this, call `winsoc_init()` before calling the function and link against the `ws2_32.lib` library.

SEE ALSO

`rpp(3B)`, `tcp(4P)`, `udp(4P)`

Chapter 8

TCL/tk Interface

The PBS Professional software includes a TCL/tk interface to PBS. Wrapped versions of many of the API calls are compiled into a special version of the TCL shell, called `pbs_tclsh`. (A special version of the tk window shell is also provided, called `pbs_wish`.) This chapter documents the TCL/tk interface to PBS.

The `pbs_tclapi` is a subset of the PBS external API wrapped in a TCL library. This functionality allows the creation of scripts that query the PBS system. Specifically, it permits the user to query the `pbs_server` about the state of PBS, jobs, queues, and nodes, and communicate with `pbs_mom` to get information about the status of running jobs, available resources on nodes, etc.

8.1 TCL/tk API Functions

A set of functions to communicate with the PBS Server and resource monitor have been added to those normally available with Tcl. All these calls will set the Tcl variable `pbs_errno` to a value to indicate if an error occurred. In all cases, the value "0" means no error. If a call to a Resource Monitor function is made, any error value will come from the system supplied `errno` variable. If the function call communicates with the PBS Server, any error value will come from the error number returned by the Server. This is the same TCL interface used by the `pbs_tclsh` and `pbs_wish` commands.

Note that the `pbs_tclapi pbsresquery` command, which calls the C API `pbs_resquery`, is deprecated. Any attempt to use it will result in a `PBSE_NOSUPPORT` error being returned.

pbs_tclapi

PBS TCL Application Programming Interface

DESCRIPTION

The `pbs_tclapi` is a subset of the PBS external API wrapped in a TCL library. This functionality allows the creation of scripts that query the PBS system. Specifically, it permits the user to query the `pbs_server` about the state of PBS, jobs, queues, and nodes, and communicate with `pbs_mom` to get information about the status of running jobs, available resources on nodes, etc.

USAGE

A set of functions to communicate with the PBS server and resource monitor have been added to those normally available with Tcl. All these calls will set the Tcl variable “`pbs_errno`” to a value to indicate if an error occurred. In all cases, the value “0” means no error. If a call to a Resource Monitor function is made, any error value will come from the system supplied `errno` variable. If the function call communicates with the PBS Server, any error value will come from the error number returned by the server. This is the same TCL interface used by the `pbs_tclsh` and `pbs_wish` commands.

`openrm host ?port?`

Creates a connection to the PBS Resource Monitor on `host` using `port` as the port number or the standard port for the resource monitor if it is not given. A connection handle is returned. If the open is successful, this will be a non-negative integer. If not, an error occurred.

`closerm connection`

The parameter `connection` is a handle to a resource monitor which was previously returned from `openrm`. This connection is closed. Nothing is returned.

`downrm connection`

Sends a command to the connected resource monitor to shutdown. Nothing is returned.

`configrm connection filename`

Sends a command to the connected resource monitor to read the configuration file given by `filename`. If this is successful, a

“0” is returned, otherwise, “-1” is returned.

addrq connection request

A resource request is sent to the connected resource monitor. If this is successful, a “0” is returned, otherwise, “-1” is returned.

getreq connection

One resource request response from the connected resource monitor is returned. If an error occurred or there are no more responses, an empty string is returned.

allreq request

A resource request is sent to all connected resource monitors. The number of streams acted upon is returned.

flushreq

All resource requests previously sent to all connected resource monitors are flushed out to the network. Nothing is returned.

activereq

The connection number of the next stream with something to read is returned. If there is nothing to read from any of the connections, a negative number is returned.

fullresp flag

Evaluates flag as a boolean value and sets the response mode used by getreq to full if flag evaluates to “true”. The full return from a resource monitor includes the original request followed by an equal sign followed by the response. The default situation is only to return the response following the equal sign. If a script needs to “see” the entire line, this function may be used.

pbsstatserv

The server is sent a status request for information about the server itself. If the request succeeds, a list with three elements is returned, otherwise an empty string is returned. The first element is the server’s name. The second is a list of attributes. The third is the “text” associated with the server (usually blank).

pbsstatjob

The server is sent a status request for information about the all jobs resident within the server. If the request succeeds, a list is returned, otherwise an empty string is returned. The list contains an entry for each job. Each element is a list with three elements. The first is the job's jobid. The second is a list of attributes. The attribute names which specify resources will have a name of the form "Resource_List:name" where "name" is the resource name. The third is the "text" associated with the job (usually blank).

pbsstatque

The server is sent a status request for information about all queues resident within the server. If the request succeeds, a list is returned, otherwise an empty string is returned. The list contains an entry for each queue. Each element is a list with three elements. This first is the queue's name. The second is a list of attributes similar to pbsstatjob. The third is the "text" associated with the queue (usually blank).

pbsstatnode

The server is sent a status request for information about all nodes defined within the server. If the request succeeds, a list is returned, otherwise an empty string is returned. The list contains an entry for each node. Each element is a list with three elements. This first is the node's name. The second is a list of attributes similar to pbsstatjob. The third is the "text" associated with the node (usually blank).

pbsselstat

The server is sent a status request for information about the all runnable jobs resident within the server. If the request succeeds, a list similar to pbsstatjob is returned, otherwise an empty string is returned.

pbsrunjob jobid ?location?

Run the job given by jobid at the location given by location. If location is not given, the default location is used. If this is successful, a "0" is returned, otherwise, "-1" is returned.

pbsasyrunjob jobid ?location?

Run the job given by jobid at the location given by location without waiting for a positive response that the job has actually

started. If location is not given, the default location is used. If this is successful, a “0” is returned, otherwise, “-1” is returned.

pbsrerunjob jobid

Re-runs the job given by jobid. If this is successful, a “0” is returned, otherwise, “-1” is returned.

pbsdeljob jobid

Delete the job given by jobid. If this is successful, a “0” is returned, otherwise, “-1” is returned.

pbsholdjob jobid

Place a hold on the job given by jobid. If this is successful, a “0” is returned, otherwise, “-1” is returned.

pbsmovejob jobid ?location?

Move the job given by jobid to the location given by location. If location is not given, the default location is used. If this is successful, a “0” is returned, otherwise, “-1” is returned.

pbsqenable queue

Set the “enabled” attribute for the queue given by queue to true. If this is successful, a “0” is returned, otherwise, “-1” is returned.

pbsqdisable queue

Set the “enabled” attribute for the queue given by queue to false. If this is successful, a “0” is returned, otherwise, “-1” is returned.

pbsqstart queue

Set the “started” attribute for the queue given by queue to true. If this is successful, a “0” is returned, otherwise, “-1” is returned.

pbsqstop queue

Set the “started” attribute for the queue given by queue to false. If this is successful, a “0” is returned, otherwise, “-1” is returned.

pbsalterjob jobid attribute_list

Alter the attributes for a job specified by `jobid`. The parameter `attribute_list` is the list of attributes to be altered. There can be more than one. Each attribute consists of a list of three elements. The first is the name, the second the resource and the third is the new value. If the alter is successful, a “0” is returned, otherwise, “-1” is returned.

`pbsresquery resource_list`

Deprecated. Obtain information about the resources specified by `resource_list`. This will be a list of strings. If the request succeeds, a list with the same number of elements as `resource_list` is returned. Each element in this list will be a list with four numbers. The numbers specify available, allocated, reserved, and down in that order.

`pbsresreserve resource_id resource_list`

Deprecated. Make (or extend) a reservation for the resources specified by `resource_list` which will be given as a list of strings. The parameter `resource_id` is a number which provides a unique identifier for a reservation being tracked by the server. If `resource_id` is given as “0”, a new reservation is created. In this case, a new identifier is generated and returned by the function. If an old identifier is used, that same number will be returned. The Tcl variable “`pbs_errno`” will be set to indicate the success or failure of the reservation.

`pbsresrelease resource_id`

Deprecated. The reservation specified by `resource_id` is released.

The two following commands are not normally used by the scheduler. They are included here because there could be a need for a scheduler to contact a server other than the one which it normally communicates with. Also, these commands are used by the Tcl tools.

`pbsconnect ?server?`

Make a connection to the named server or the default server if a parameter is not given. Only one connection to a server is allowed at any one time.

`pbsdisconnect`

Disconnect from the currently connected server.

The above Tcl functions use PBS interface library calls for communication with the server and the PBS resource monitor library to communicate with pbs_mom.

`datetime ?day? ?time?`

The number of arguments used determine the type of date to be calculated. With no arguments, the current POSIX date is returned. This is an integer in seconds.

With one argument there are two possible formats. The first is a 12 (or more) character string specifying a complete date in the following format:

YYMMDDhhmmss

All characters must be digits. The year (YY) is given by the first two (or more) characters and is the number of years since 1900. The month (MM) is the number of the month [01-12]. The day (DD) is the day of the month [01-32]. The hour (hh) is the hour of the day [00-23]. The minute (mm) is minutes after the hour [00-59]. The second (ss) is seconds after the minute [00-59]. The POSIX date for the given date/time is returned.

The second option with one argument is a relative time. The format for this is

HH:MM:SS

With hours (HH), minutes (MM) and seconds (SS) being separated by colons “:”. The number returned in this case will be the number of seconds in the interval specified, not an absolute POSIX date.

With two arguments a relative date is calculated. The first argument specifies a day of the week and must be one of the following strings: “Sun”, “Mon”, “Tue”, “Wed”, “Thr”, “Fri”, or “Sat”. The second argument is a relative time as given above. The POSIX date calculated will be the day of the week given which follows the current day, and the time given in the second argument. For example, if the current day was Monday, and the two arguments were “Fri” and “04:30:00”, the date calculated would be the POSIX date for the Friday following the current Monday, at four-thirty in the morning. If the day specified and the current day are the same, the current day is used, not the day one week later.

strftime format time

This function calls the POSIX function `strftime()`. It requires two arguments. The first is a format string. The format conventions are the same as those for the POSIX function `strftime()`. The second argument is POSIX calendar time in seconds as returned by `datetime`. It returns a string based on the format given. This gives the ability to extract information about a time, or format it for printing.

logmsg tag message

This function calls the internal PBS function `log_err()`. It will cause a log message to be written to the scheduler's log file. The `tag` specifies a function name or other word used to identify the area where the message is generated. The message is the string to be logged.

SEE ALSO

`pbs_tclsh(8B)`, `pbs_wish(8B)`, `pbs_mom(8B)`, `pbs_server(8B)`,
`pbs_sched(8B)`

Chapter 9

Hooks

This chapter describes the PBS hook APIs. For more information on hooks, see the PBS Professional Administrator's Guide.

9.1 Introduction

A hook is a block of Python code that is triggered in response to queuing a job, modifying a job, moving a job, running a job, or submitting a PBS reservation. Each hook can *accept* (allow) or *reject* (prevent) the action that triggers it. The hook can modify the input parameters given for the action. The hook can also make calls to functions external to PBS. PBS provides an interface for use by hooks. This interface allows hooks to read and/or modify things such as job and server attributes, the server, queues, and the event that triggered the hook.

The Administrator creates any desired hooks. No special configuration of PBS is required in order to use hooks.

This chapter contains the following man pages:

- `pbs_module(7B)`
- `pbs_stathook(3B)`
- `pbs_hook_attributes(7B)`

See the following additional man pages:

- `qmgr(1B)`
- `qsub(1B)`
- `qmove(1B)`
- `qalter(1B)`
- `pbs_rsub(1B)`
- `pbs_manager(3B)`

9.2 How Hooks Work

9.2.1 Hook Contents and Permissions

A hook contains a Python script. The script is evaluated by a Python 2.5 or later interpreter, embedded in PBS.

Hooks have a default UNIX umask of 022. File permissions are inherited from the current working directory of the hook script.

9.2.2 Accepting and Rejecting Actions

The hook script always accepts the current event request action unless an unhandled exception occurs in the script, a hook alarm timeout is triggered or there's an explicit call to `“pbs.event().reject()”`.

9.2.3 Exceptions

A hook script can catch an exception and evaluate whether or not to accept or reject the event action. In this example, while referencing the non-existent attribute `pbs.event().job.interactive`, an exception is triggered, but the event action is still accepted:

```
...
try:
    e = pbs.event()
    if e.job.interactive:
        e.reject("Interactive jobs not allowed")
except SystemExit:
    pass
except:
    e.accept()
```

9.2.4 Unsupported Interfaces and Uses

Site hooks which read, write, close, or alter `stdin`, `stdout`, or `stderr`, are not supported. Hooks which use any interfaces other than those described are unsupported.

9.3 Interface to Hooks

Two PBS APIs are used with hooks. These are `pbs_manager()` and `pbs_stathook()`. The `pbs` module provides a Python interface to PBS.

9.3.1 The `pbs` Module

Hooks have access to a special module called “`pbs`”, which contains functions that perform PBS-related actions. This module must be explicitly loaded by the hook writer via the call “`import pbs`”.

The *pbs module* provides an interface to PBS and the hook environment. The interface is made up of Python objects, which have attributes and methods. You can operate on these objects using Python code.

For a description of the `pbs` module, see section 6.9, “Interface to Hooks”, on page 402 of the PBS Professional Administrator’s Guide.

9.3.2 The `pbs_manager()` API

The `pbs_manager()` API is described in “`pbs_manager`” on page 42. The elements related to hooks are repeated here:

The `pbs_manager()` API contains an `obj_name` called “hook” defined as `MGR_OBJ_HOOK`.

To run, hooks require root privilege on UNIX, and local Administrators privilege on Windows. Hooks run only on the server host.

The `pbs_manager()` API contains the following hook commands, which operate only on hook objects:

MGR_CMD_IMPORT

This command is used for loading the hook script contents into a hook.

MGR_CMD_EXPORT

This command is used for dumping to a file the contents of a hook script.

The parameters to `MGR_CMD_IMPORT` and `MGR_CMD_EXPORT` are specified via the `attrib` parameter of `pbs_manager()`. The `attrib` parameter is a “struct `attropl`” defined in `pbs_if1.h` as:

```
struct attropl {
    char *name;
    char *resource;
    char *value;
    enum batch_op op;
    struct attropl *next;
};
```

The `attrib` list is terminated by the first entry where `next` is null.

For `MGR_CMD_IMPORT`, specify `attropl 'name'` as “content-type”, “content-encoding”, and “input-file” along with the corresponding `'value'` and an `'op'` of SET.

For `MGR_CMD_EXPORT`, specify the `attropl 'name'` as “content-type”, “content-encoding”, and “output-file” along with the corresponding `'value'`, and an `'op'` of SET.

9.3.2.1 Examples of Using `pbs_manager()`

Example 9-1: The following:

```
# qmgr -c 'import hook hook1 application/x-python base64 hello.py.b64'
```

is programmatically equivalent to:

```
static struct attrop1 imp_attribs[] = {
    { "content-type",
      (char *)0,
      "application/x-python",
      SET,
      (struct attrop1 *)&imp_attribs[1]
    },
    { "content-encoding",
      (char *)0,
      "base64",
      SET,
      (struct attrop1 *)&imp_attribs[2]},
    { "input-file",
      (char *)0,
      "hello.py.b64",
      SET,
      (struct attrop1 *)0
    }
};

pbs_manager(con, MGR_CMD_IMPORT, MGR_OBJ_HOOK, "hook1", &imp_attribs[0],
            NULL);
```

Example 9-2: The following:

```
# qmgr -c 'export hook hook1 application/x-python default hello.py'
```

is programmatically equivalent to:

```
static struct attrtpl exp_attribs[] = {
    { "content-type",
      (char *)0,
      "application/x-python",
      SET,
      (struct attrtpl *)&exp_attribs[1]},
    { "content-encoding",
      (char *)0,
      "default",
      SET,
      (struct attrtpl *)&exp_attribs[2]},
    { "output-file",
      (char *)0,
      "hello.py",
      SET,
      (struct attrtpl *)0
    }
};

pbs_manager(con, MGR_CMD_EXPORT, MGR_OBJ_HOOK, "hook1", &exp_attribs[0],
            NULL);
```

9.3.3 The pbs_stathook() API

The PBS API called “pbs_stathook()” is used to get attributes and values for site hooks.

The prototype for pbs_stathook() is as follows:

```
struct batch_status *pbs_stathook(int connect, char *hook_name, struct
    attrl *attrib, char *extend)
```

The call to pbs_stathook() causes a PBS_BATCH_StatusHook request to be sent to the server. In reply, the PBS server returns a batch reply status of object type MGR_OBJECT_HOOK listing the attributes and values that were requested relating to a particular hook or all hooks of type HOOK_SITE.

9.3.3.1 Example of Using `pbs_stathook()`

To list all site hooks using `qmgr`:

```
qmgr -c "list hook"
```

To list all site hooks using the `pbs_stathook()` API:

```
pbs_stathook()
```

The result is the same. For example, if there are two site hooks, `c3` and `c36`:

```
Hook c3
```

```
type = site
enabled = true
event = queuejob, modifyjob
user = pbsadmin
alarm = 30
order = 1
```

```
Hook c36
```

```
type = site
enabled = true
event = resvsub
user = pbsadmin
alarm = 30
order = 1
```

9.4 The pbs Module

pbs_module

DESCRIPTION

The `pbs` module provides an interface to PBS and the hook environment. The interface is made up of Python objects, which have attributes and methods. You can operate on these objects using Python code. For a description of each object, see the *PBS Professional Administrator's Guide*.

PBS MODULE OBJECTS

`pbs.acl`

Represents a PBS ACL type.

`pbs.args`

Represents a space-separated list of PBS arguments to commands like `qsub`, `qdel`.

`pbs.BadAttributeValueError`

Raised when setting the attribute value of a `pbs.*` object to an invalid value.

`pbs.BadAttributeValueTypeError`

Raised when setting the attribute value of a `pbs.*` object to an invalid value type.

`pbs.BadResourceValueError`

Raised when setting the resource value of a `pbs.*` object to an invalid value.

`pbs.BadResourceValueTypeError`

Raised when setting the resource value of a `pbs.*` object to an invalid value type.

`pbs.checkpoint`

Represents a job's checkpoint attribute.

`pbs.depend`

Represents a job's depend attribute.

- `pbs.duration`
Represents a time interval.
- `pbs.email_list`
Represents the set of users to whom mail may be sent.
- `pbs.event`
Represents a PBS event.
- `pbs.EventIncompatibleError`
Raised when referencing a non-existent attribute in `pbs.event()`.
- `pbs.exec_host`
Represents a job's `exec_host` attribute.
- `pbs.exec_vnode`
Represents a job's `exec_vnode` attribute.
- `pbs.group_list`
Represents a list of group names.
- `pbs.hold_types`
Represents a job's `Hold_Types` attribute.
- `pbs.job`
Represents a PBS job.
- `pbs.job_sort_formula`
Represents the server's `job_sort_formula` attribute.
- `pbs.JOB_STATE_BEGUN`
Represents the job array state of having started.
- `pbs.JOB_STATE_EXITING`
Represents the job state of exiting.
- `pbs.JOB_STATE_EXPIRED`
Represents the subjob state of expiring.
- `pbs.JOB_STATE_FINISHED`
Represents the job state of finished.

`pbs.JOB_STATE_HELD`

Represents the job state of held.

`pbs.JOB_STATE_MOVED`

Represents the job state of moved.

`pbs.JOB_STATE_QUEUED`

Represents the job state of queued.

`pbs.JOB_STATE_RUNNING`

Represents the job state of running.

`pbs.JOB_STATE_SUSPEND`

Represents the job state of suspended.

`pbs.JOB_STATE_SUSPEND_USERACTIVE`

Represents the job state of suspended due to user activity.

`pbs.JOB_STATE_TRANSIT`

Represents the job state of transiting.

`pbs.JOB_STATE_WAITING`

Represents the job state of waiting.

`pbs.join_path`

Represents a job's `Join_Path` attribute.

`pbs.keep_files`

Represents a job's `Keep_Files` attribute.

`pbs.license_count`

Represents a set of licensing-related counters.

`pbs.logjobmsg()`

Message logging method taking job ID and message as arguments.

`pbs.logmsg()`

Message logging method taking log level and message as arguments.

`pbs.LOG_DEBUG`

Log level 004.

pbs.LOG_ERROR
Log level 004.

pbs.LOG_WARNING
Log level 004.

pbs.mail_points
Represents a job's Mail_Points attribute.

pbs.MODIFYJOB
The modifyjob event type.

pbs.MOVEJOB
The movejob event type.

pbs.ND_BUSY
Represents busy vnode state.

pbs.ND_DEFAULT_EXCL
Represents default_excl sharing vnode attribute value

pbs.ND_DEFAULT_SHARED
Represents default_shared sharing vnode attribute value.

pbs.ND_DOWN
Represents down vnode state

pbs.ND_FORCE_EXCL
Represents force_excl sharing vnode attribute value.

pbs.ND_FREE
Represents free vnode state.

pbs.ND_GLOBUS
Represents globus vnode ntype.

pbs.ND_IGNORE_EXCL
Represents ignore_excl sharing vnode attribute value.

pbs.ND_JOBBUSY

Represents job-busy vnode state.

`pbs.ND_JOB_EXCLUSIVE`
Represents job-exclusive vnode state.

`pbs.ND_OFFLINE`
Represents offline vnode state.

`pbs.ND_PBS`
Represents pbs vnode ntype.

`pbs.ND_PROV`
Represents provisioning vnode state.

`pbs.ND_STALE`
Represents stale vnode state.

`pbs.ND_STATE_UNKNOWN`
Represents state-unknown, down vnode state.

`pbs.ND_WAIT_PROV`
Represents wait-provisioning vnode state.

`pbs.node_group_key`
Represents the server or queue `node_group_key` attribute.

`pbs.path_list`
Represents a list of pathnames.

`pbs.place`
Represents the place job submission specification.

`pbs.QTYPE_EXECUTION`
The execution queue type.

`pbs.QTYPE_ROUTE`
The route queue type.

`pbs.queue`
Represents a PBS queue.

`pbs.QUEUEJOB`

The queuejob event type.

`pbs.range`

Represents a range of numbers referring to array indices.

`pbs.resv`

Represents a PBS reservation.

`pbs.RESVSUB`

The resvsub event type.

`pbs.RESV_STATE_BEING_DELETED`

Represents the reservation state `RESV_BEING_DELETED`.

`pbs.RESV_STATE_CONFIRMED`

Represents the reservation state `RESV_CONFIRMED`.

`pbs.RESV_STATE_DEGRADED`

Represents the reservation state `RESV_DEGRADED`.

`pbs.RESV_STATE_DELETED`

Represents the reservation state `RESV_DELETED`.

`pbs.RESV_STATE_DELETING_JOBS`

Represents the reservation state `RESV_DELETING_JOBS`.

`pbs.RESV_STATE_FINISHED`

Represents the reservation state `RESV_FINISHED`.

`pbs.RESV_STATE_NONE`

Represents the reservation state `RESV_NONE`.

`pbs.RESV_STATE_RUNNING`

Represents the reservation state `RESV_RUNNING`.

`pbs.RESV_STATE_TIME_TO_RUN`

Represents the reservation state `RESV_TIME_TO_RUN`.

`pbs.RESV_STATE_UNCONFIRMED`

Represents the reservation state `RESV_UNCONFIRMED`.

`pbs.RESV_STATE_WAIT`

Represents the reservation state RESV_WAIT.

`pbs.route_destinations`

Represents a queue's `route_destinations` attribute.

`pbs.RUNJOB`

The `runjob` event type.

`pbs.select`

Represents the select job submission specification.

`pbs.server`

Represents the local PBS server.

`pbs.size`

Represents a PBS size type.

`pbs.software`

Represents a site-dependent software specification resource.

`pbs.staging_list`

Represents a list of file stagein or stageout parameters.

`pbs.state_count`

Represents a set of job-related state counters.

`pbs.SV_STATE_ACTIVE`

Represents the server state "Scheduling".

`pbs.SV_STATE_HOT`

Represents the server state "Hot_Start".

`pbs.SV_STATE_IDLE`

Represents the server state "Idle".

`pbs.SV_STATE_SHUTDEL`

Represents the server state "Terminating, Delayed".

`pbs.SV_STATE_SHUTIMM`

Represents the server state "Terminating".

`pbs.SV_STATE_SHUTSIG`

Represents the server state “Terminating”, when a signal has been caught.

`pbs.UnsetAttributeNameError`

Raised when referencing a non-existent name of a `pbs.*` object.

`pbs.UnsetResourceNameError`

Raised when referencing a non-existent name of a `pbs.*` object.

`pbs.user_list`

Represents a list of user names.

`pbs.vchunk`

Represents a resource chunk assigned to a job.

`pbs.version`

Represents PBS version information.

`pbs.vnode`

Represents a PBS vnode.

`SystemExit`

Raised when accepting or rejecting an action.

Resources

Hooks can read `Server`, `Queue`, `reservation`, `vnode`, or `job` resources.

Hooks can modify

- The resources requested by a job

- The resource arguments to `pbs_rsub`

- The state `vnode` attribute

Custom and built-in PBS resources are represented in Python dictionaries, where the resource names are the dictionary keys. Built-in resources are listed in `pbs_resources(7B)`. You reference a resource through either the `Server` or the event that triggered the hook, for example:

```
pbs.server().resources_available[“< resource name>”]
pbs.event().job.Resource_List[“< resource name>”]
```

The resource name must be in quotes. Example: Get the number of CPUs:

```
ncpus = Resource_List["ncpus"]
```

An instance R of a job resource can be set as follows:

```
R["<resource name>"] = <resource value>
```

For example:

```
pbs.event().job().Resource_List["mem"] = 8gb
```

EXCEPTIONS

`pbs.BadAttributeValueError`

Raised when setting the attribute value of a `pbs.*` object to an invalid value.

`pbs.BadAttributeValueTypeError`

Raised when setting the attribute value of a `pbs.*` object to an invalid value type.

`pbs.BadResourceValueError`

Raised when setting the resource value of a `pbs.*` object to an invalid value.

`pbs.BadResourceValueTypeError`

Raised when setting the resource value of a `pbs.*` object to an invalid value type.

`pbs.EventIncompatibleError`

Raised when referencing a non-existent attribute in `pbs.event()`.

`pbs.UnsetAttributeNameError`

Raised when referencing a non-existent name of a `pbs.*` object.

`pbs.UnsetResourceNameError`

Raised when referencing a non-existent name of a `pbs.*` object.

`SystemExit`

Raised when accepting or rejecting an action.

If a hook encounters an unhandled exception:

PBS rejects the corresponding action, and an error message is

printed to stderr.
A message is printed to the daemon log.

SEE ALSO

The PBS Professional Administrator's Guide,
pbs_hook_attributes(7B), pbs_resources(7B), qmgr(1B)

9.5 The pbs_stathook() API

pbs_stathook(3B)

SYNOPSIS

```
#include <pbs_error.h>
#include <pbs_ifl.h>

struct batch_status *pbs_stathook(int connect, char *id,
struct attrl *attrib, char *extend)

void pbs_statfree(struct batch_status *psj)
```

DESCRIPTION

Issue a batch request to obtain the status of a specified site hook or a set of site hooks at the current server.

A Status Hook batch request is generated and sent to the server over the connection specified by `connect` which is the return value of `pbs_connect()`.

This API can be executed only by root on the local server host.

The parameter, `id`, may be either a hook name or the null string. If `id` specifies a name, the attribute-value list for that hook is returned. If `id` is a null string or a null pointer, the status of a all hooks at the current server is returned.

The parameter, `attrib`, is a pointer to an `attrl` structure which is defined in `pbs_ifl.h` as:

```
struct attrl {
    struct attrl *next;
    char *name;
    char *resource;
```

```

    char    *value;
};

```

The attrib list is terminated by the first entry where next is a null pointer.

If an attrib list is given, then only the attributes in the list are returned by the server. Otherwise, all the attributes of a hook are returned.

The resource member is only used if the name member is ATTR_1, otherwise it should be a pointer to a null string.

The value member should always be a pointer to a null string.

The parameter, extend, is reserved for implementation defined extensions.

The return value is a pointer to a list of batch_status structures or the null pointer if no site hooks can be queried for status. The batch_status structure is defined in pbs_ifl.h as

```

struct batch_status {
    struct batch_status *next;
    char    *name;
    struct attrl    *attribs;
    char    *text;
}

```

It is up to the user to free the structure when no longer needed, by calling pbs_statfree().

SEE ALSO

pbs_hook_attributes(7B), pbs_connect(3B)

DIAGNOSTICS

When the batch request generated by the pbs_stathook() function has been completed successfully and the status of each site hook has been returned by the batch server, the routine will return a pointer to the list of batch_status structures. If no site hooks were available to query or an error occurred, a null pointer is returned. The global integer pbs_errno should be examined to determine the cause.

Chapter 10

HPC Basic Profile

This chapter describes using PBS for HPC Basic Profile jobs. For more information on HPC Basic Profile, see Chapter 7, "Metascheduling Using HPC Basic Profile", on page 481 in the PBS Professional Administrator's Guide.

10.1 Introduction

PBS Professional can schedule and manage jobs on one or more HPC Basic Profile Servers using the Grid Forum OGSA HPC Basic Profile web services standard.

For definitions, see section 7.1, "Definitions", on page 481 of the PBS Professional Administrator's Guide.

10.2 How PBS Works With HPC Basic Profile

10.2.1 Information Available From HPCBP Nodes

The only information available about the HPCBP nodes is that which is supplied by the HPC Basic Profile Server to the HPCBP MOM.

10.2.2 Translating Jobs for HPCBP

10.2.2.1 Translating Job Attributes for HPCBP Jobs

The HPCBP MOM converts the job's attributes to equivalent JSDL elements, and passes the resulting JSDL document to the HPC Basic Profile Server.

If the select statement includes HPCBP hostnames, those names are passed using the *<CandidateHosts>* JSDL element.

The aggregate number of CPUs is passed using the *<TotalCPUCount>* JSDL element.

See the following table for a mapping of job attributes to JSDL elements:

Table 10-1: How PBS Job Attributes Are Translated Into JSDL

PBS Attribute/Directive	JSDL element
-l select = arch = <value>	<i><OperatingSystemVersion></i> or <i><OperatingSystemName></i>
-l cput = <time>	<i><TotalCPUTime></i>
-l file = <size>	<i><IndividualDiskSpace></i>
-l select = N:host = <value> Or -l select = N:vnode=<value>	<i><HostName></i>
-l select = N:mem = <size>	<i><TotalPhysicalMemory></i>
-l select = N :ncpus = <value> :N:mpiprocs = <value>	<i><TotalCPUCount></i>
-l pcput = <time>	<i><IndividualCPUTime></i>
-l pmem = <size>	<i><IndividualPhysicalMemory></i>
-l pvmem = <size>	<i><IndividualVirtualMemory></i>
-l vmem = <size>	<i><TotalVirtualMemory></i>
-N Job_Name	<i><JobName></i>
-w stagein = <directive>	<i><DataStaging></i> + <i><Source></i>
-w stageout = <directive>	<i><DataStaging></i> + <i><Target></i>
-o <hostname> : filepath	<i><DataStaging></i> + <i><Target></i>
-e <hostname> : filepath	<i><DataStaging></i> + <i><Target></i>
-j [o][e]	<i><DataStaging></i> + <i><Target></i>
-k [o][e]	----

Table 10-1: How PBS Job Attributes Are Translated Into JSDL

PBS Attribute/Directive	JSDL element
Variable_List	<Environment>
-l ompthreads	<Environment>

10.2.2.2 Translating arch Values for HPCBP Jobs

The following table lists arch values and their *OperatingSystemName* and *OperatingSystemVersion* JSDL equivalents:

Table 10-2: Architectures Used in HPCBP Jobs

arch	OperatingSystem Name	OperatingSystemVersion
<i>aix4</i>	<i>AIX</i>	----
<i>hpux10</i>	<i>HPUX</i>	----
<i>hpux11</i>	<i>HPUX</i>	----
<i>irix6</i>	<i>IRIX</i>	----
<i>linux</i>	<i>LINUX</i>	----
<i>solaris7</i>	<i>Solaris</i>	----
<i>Digitalunix</i>	<i>Tru64_UNIX</i>	----
<i>windows_2003_server</i>	<i>other</i>	<i>Windows_2003_server</i>
<i>windows_2008_server</i>	<i>other</i>	<i>Windows_2008_server</i>
<i>other</i>	<i>other</i>	----
<i>macos</i>	<i>MACOS</i>	----
<i>attunix</i>	<i>ATTUNIX</i>	----
<i>dgux</i>	<i>DGUX</i>	----
<i>decnt</i>	<i>DECNT</i>	----

Table 10-2: Architectures Used in HPCBP Jobs

arch	OperatingSystem Name	OperatingSystemVersion
<i>OpenMVS</i>	<i>OpenMVS</i>	----
<i>MVS</i>	<i>MVS</i>	----
<i>OS400</i>	<i>OS400</i>	----
<i>OS_2</i>	<i>OS_2</i>	----
<i>JavaVM</i>	<i>JavaVM</i>	----
<i>WINNT</i>	<i>WINNT</i>	----
<i>WINCE</i>	<i>WINCE</i>	----
<i>NCR3000</i>	<i>NCR3000</i>	----
<i>NetWare</i>	<i>NetWare</i>	----
<i>OSF</i>	<i>OSF</i>	----
<i>DC_OS</i>	<i>DC_OS</i>	----
<i>Reliant_UNIX</i>	<i>Reliant_UNIX</i>	----
<i>SCO_UnixWare</i>	<i>SCO_UnixWare</i>	----
<i>SCO_OpenServer</i>	<i>SCO_OpenServer</i>	----
<i>Sequent</i>	<i>Sequent</i>	----
<i>SunOS</i>	<i>SunOS</i>	----
<i>U6000</i>	<i>U6000</i>	----
<i>ASERIES</i>	<i>ASERIES</i>	----
<i>TandemNSK</i>	<i>TandemNSK</i>	----
<i>TandemNT</i>	<i>TandemNT</i>	----
<i>BS2000</i>	<i>BS2000</i>	----
<i>Lynx</i>	<i>Lynx</i>	----

Table 10-2: Architectures Used in HPCBP Jobs

arch	OperatingSystem Name	OperatingSystemVersion
<i>XENIX</i>	<i>XENIX</i>	----
<i>VM</i>	<i>VM</i>	----
<i>Interactive_UNIX</i>	<i>Interactive_UNIX</i>	----
<i>BSDUNIX</i>	<i>BSDUNIX</i>	----
<i>FreeBSD</i>	<i>FreeBSD</i>	----
<i>NetBSD</i>	<i>NetBSD</i>	----
<i>GNU_Hurd</i>	<i>GNU_Hurd</i>	----
<i>OS9</i>	<i>OS9</i>	----
<i>MACH_Kernel</i>	<i>MACH_Kernel</i>	----
<i>Inferno</i>	<i>Inferno</i>	----
<i>QNX</i>	<i>QNX</i>	----
<i>EPOC</i>	<i>EPOC</i>	----
<i>IxWorks</i>	<i>IxWorks</i>	----
<i>VxWorks</i>	<i>VxWorks</i>	----
<i>MiNT</i>	<i>MiNT</i>	----
<i>BeOS</i>	<i>BeOS</i>	----
<i>HP_MPE</i>	<i>HP_MPE</i>	----
<i>NextStep</i>	<i>NextStep</i>	----
<i>PalmPilot</i>	<i>PalmPilot</i>	----
<i>Rhapsody</i>	<i>Rhapsody</i>	----
<i>Windows_2000</i>	<i>Windows_2000</i>	----
<i>Dedicated</i>	<i>Dedicated</i>	----

Table 10-2: Architectures Used in HPCBP Jobs

arch	OperatingSystem Name	OperatingSystemVersion
<i>OS_390</i>	<i>OS_390</i>	----
<i>VSE</i>	<i>VSE</i>	----
<i>TPF</i>	<i>TPF</i>	----
<i>Windows_R_Me</i>	<i>Windows_R_Me</i>	----
<i>Caldera_Open_UNIX</i>	<i>Caldera_Open_UNI X</i>	----
<i>OpenBSD</i>	<i>OpenBSD</i>	----
<i>Not_Applicable</i>	<i>Not_Applicable</i>	----
<i>Windows_XP</i>	<i>Windows_XP</i>	----
<i>z_OS</i>	<i>z_OS</i>	----
<i><any other string></i>	<i>other</i>	<i><any other string></i>

10.3 Examples

Example 10-1: PBS job request:

```
qsub -N job1 -- /bin/executable -a -b -cd
```

JSDL document:

```
<?xml version="1.0" encoding="utf-8"?>
<jsd1:JobDefinition xmlns:jsdl="http://schemas.ggf.org/jsdl/2005/11/jsdl"
  xmlns:jsdl-hpcpa="http://schemas.ggf.org/jsdl/2006/07/jsdl-hpcpa"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" >
  <jsd1:JobDescription>
    <jsd1:JobIdentification>
      <jsd1:JobName>job1</jsdl:JobName>
      <jsd1:JobProject> PBS-TO-HPCBP v1</jsdl:Job-
        Project>
    </jsdl:JobIdentification>
    <jsd1:Application>
      <jsd1-hpcpa:HPCProfileApplication name="PBS-
        toHPCBP">
      <jsd1-hpcpa:Executable>/bin/execu-
        table</jsdl-hpcpa:Executable>
      <jsd1-hpcpa:Argument>-a</jsdl-hpcpa:Argument>
      <jsd1-hpcpa:Argument>b</jsdl-hpcpa:Argument>
      <jsd1-hpcpa:Argument>-cd</jsdl-hpcpa:Argument>
      </jsdl-hpcpa:HPCProfileApplication>
    </jsdl:Application>
  </jsdl:JobDescription>
</jsdl:JobDefinition>
```

Example 10-2: PBS job request:

```
qsub -N job2 -l select=2:n_cpus=2:mpiprocs=2 -- /bin/executable a b c
```

JSDL document:

```
<?xml version="1.0" encoding="utf-8"?>
<jsdsl:JobDefinition xmlns:jsdl="http://schemas.ggf.org/jsdl/2005/11/jsdl"
  xmlns:jsdl-hpcpa="http://schemas.ggf.org/jsdl/2006/07/jsdl-hpcpa"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" >
  <jsdsl:JobDescription>
    <jsdsl:JobIdentification>
      <jsdsl:JobName>job2</jsdl:JobName>
      <jsdsl:JobProject> PBS-TO-HPCBP v1</jsdl:Job-
        Project>
    </jsdl:JobIdentification>
    <jsdsl:Application>
      <jsdsl-hpcpa:HPCProfileApplication name="PBS-
        toHPCBP" >
      <jsdsl-hpcpa:Executable>/bin/execu-
        table</jsdl-hpcpa:Executable>
      <jsdsl-hpcpa:Argument>a</jsdl-hpcpa:Argument>
      <jsdsl-hpcpa:Argument>b</jsdl-hpcpa:Argument>
      <jsdsl-hpcpa:Argument>c</jsdl-hpcpa:Argument>
      </jsdl-hpcpa:HPCProfileApplication>
    </jsdl:Application>
    <jsdsl:Resources>
      <jsdsl:TotalCPUCount>
        <jsdsl:UpperBoundedRange>4</jsdl:Upper-
          BoundedRange>
        <jsdsl:LowerBoundedRange>4</jsdl:Lower-
          BoundedRange>
      </jsdl:TotalCPUCount>
    </jsdl:Resources>
  </jsdl:JobDescription>
</jsdl:JobDefinition>
```

Example 10-3: PBS job request:

```
qsub -N job3 -l select= 4 :ncpus=4 :arch=windows_2008_server
      :host=winhost0+2:ncpus=2: host=winhost1 -- /bin/executable a b
```

JSDL document:

```
<?xml version="1.0" encoding="utf-8"?>
<jsdsl:JobDefinition      xmlns:jsdsl="http://schemas.ggf.org/jsdsl/2005/11/
jsdsl" xmlns:jsdsl-hpcpa="http://schemas.ggf.org/jsdsl/2006/07/jsdsl-
hpcpa" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" >
  <jsdsl:JobDescription>
    <jsdsl:JobIdentification>
      <jsdsl:JobName>job3</jsdsl:JobName>
      <jsdsl:JobProject> PBS-TO-HPCBP v1</jsdsl:Job-
        Project>
    </jsdsl:JobIdentification>
    <jsdsl:Application>
      <jsdsl-hpcpa:HPCProfileApplication name="PBSto-
        HPCBP" >
      <jsdsl-hpcpa:Executable>/bin/execu-
        table</jsdsl-hpcpa:Executable>
      <jsdsl-hpcpa:Argument>a</jsdsl-hpcpa:Argument>
      <jsdsl-hpcpa:Argument>b</jsdsl-hpcpa:Argument>
      </jsdsl-hpcpa:HPCProfileApplication>
    </jsdsl:Application>
    <jsdsl:Resources>
      <jsdsl:CandidateHosts>
        <jsdsl:HostName>winhost0</jsdsl:HostName>
        <jsdsl:HostName>winhost1</jsdsl:HostName>
      </jsdsl:CandidateHosts>
      <jsdsl:OperatingSystem>
        <jsdsl:OperatingSystemType>
          <jsdsl:OperatingSystemName>oth-
            er</jsdsl:OperatingSystemName>
        </jsdsl:OperatingSystemType>
        <jsdsl:OperatingSystemVersion>Windows_2008_-
          Server</jsdsl:OperatingSystemVersion>
      </jsdsl:OperatingSystem>
```

```
<jSDL:TotalCPUCount>
  <jSDL:UpperBoundedRange>20</jSDL:Upper-
    BoundedRange>
  <jSDL:LowerBoundedRange>20</jSDL:Lower-
    BoundedRange>
</jSDL:TotalCPUCount>
</jSDL:Resources>
</jSDL:JobDescription>
</jSDL:JobDefinition>
```

Example 10-4: PBS job request:

```
qsub -N job4 -l select=4:n_cpus=2:mpi_procs=2 -- /bin/executable a b
```

JSDL document:

```
<?xml version="1.0" encoding="utf-8"?>
<jsdsl:JobDefinition      xmlns:jsdl="http://schemas.ggf.org/jsdl/2005/11/
  jsdl" xmlns:jsdl-hpcpa="http://schemas.ggf.org/jsdl/2006/07/jsdl-
  hpcpa" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" >
  <jsdsl:JobDescription>
    <jsdsl:JobIdentification>
      <jsdsl:JobName>job4</jsdl:JobName>
      <jsdsl:JobProject> PBS-TO-HPCBP v1</jsdl:Job-
        Project>
    </jsdl:JobIdentification>
    <jsdsl:Application>
      <jsdsl-hpcpa:HPCProfileApplication name="PBSto-
        HPCBP" >
        <jsdsl-hpcpa:Executable>/bin/execu-
          table</jsdl-hpcpa:Executable>
        <jsdsl-hpcpa:Argument>a</jsdl-hpcpa:Argu-
          ment>
        <jsdsl-hpcpa:Argument>b</jsdl-hpcpa:Argu-
          ment>
      </jsdl-hpcpa:HPCProfileApplication>
    </jsdl:Application>
    <jsdsl:Resources>
      <jsdsl:TotalCPUCount>
        <jsdsl:UpperBoundedRange>8</jsdl:UpperBound-
          edRange>
        <jsdsl:LowerBoundedRange>8</jsdl:LowerBound-
          edRange>
      </jsdl:TotalCPUCount>
    </jsdl:Resources>
  </jsdl:JobDescription>
</jsdl:JobDefinition>
```

Example 10-5: Job that specifies all the PBS job attributes:

```
qsub -N job5 -l select=4:ncpus=2:mpiprocs=2:arch=windows_2008_server:
      host=winhost0:ompthreads=2 -l cput=00:01:20 -l file=2kb -l
      pcpur=00:00:10 -l pmem=30kb -l pvmem=30kb -l vmem=240kb -- /bin/execut-
      able a b
```

JSDL document:

```
<?xml version="1.0" encoding="utf-8"?>
<jSDL:JobDefinition      xmlns:jSDL="http://schemas.ggf.org/jSDL/2005/11/
jSDL" xmlns:jSDL-hPCPA="http://schemas.ggf.org/jSDL/2006/07/jSDL-
hPCPA" xmlns:xSI="http://www.w3.org/2001/XMLSchema-instance" >
  <jSDL:JobDescription>
    <jSDL:JobIdentification>
      <jSDL:JobName>job5</jSDL:JobName>
      <jSDL:JobProject> PBS-TO-HPCBP v1</jSDL:Job-
        Project>
    </jSDL:JobIdentification>
    <jSDL:Application>
      <jSDL-hPCPA:HPCProfileApplication name="PBSto-
        HPCBP" >
        <jSDL-hPCPA:Executable>/bin/execu-
          table</jSDL-hPCPA:Executable>
        <jSDL-hPCPA:Argument>a</jSDL-hPCPA:Argu-
          ment>
        <jSDL-hPCPA:Argument>b</jSDL-hPCPA:Argu-
          ment>
        <jSDL-hPCPA:Environment name="OMP_NUM_-
          THREADS"> 2</jSDL-hPCPA:Environment>
      </jSDL-hPCPA:HPCProfileApplication>
    </jSDL:Application>
  <jSDL:Resources>
    <jSDL:CandidateHosts>
      <jSDL:HostName>winhost0</jSDL:HostName>
    </jSDL:CandidateHosts>
    <jSDL:OperatingSystem>
      <jSDL:OperatingSystemType>
        <jSDL:OperatingSystemName>oth-
```

```

er</jsdl:OperatingSystemName>
</jsdl:OperatingSystemType>
<jsdl:OperatingSystemVersion>Windows_2008_
  Server</jsdl:OperatingSystemVersion>
</jsdl:OperatingSystem>
<jsdl:TotalCPUCount>
  <jsdl:UpperBoundedRange>8</jsdl:UpperBound-
    edRange>
  <jsdl:LowerBoundedRange>8</jsdl:LowerBound-
    edRange>
</jsdl:TotalCPUCount>
<jsdl:TotalCPUTime>
  <jsdl:UpperBoundedRange>80.0</jsdl:Upper-
    BoundedRange>
</jsdl:TotalCPUTime>
<jsdl:IndividualDiskSpace>
  <jsdl:UpperBoundedRange>2048.0</jsdl:Upper-
    BoundedRange>
</jsdl:IndividualDiskspace>
<jsdl:IndividualCPUTime>
  <jsdl:UpperBoundedRange>10.0</jsdl:Upper-
    BoundedRange>
</jsdl:IndividualCPUTime>
<jsdl:IndividualPhysicalMemory>
  <jsdl:UpperBoundedRange>30720.0</jsdl:Up-
    perBoundedRange>
</jsdl:IndividualPhysicalMemory>
<jsdl:IndividualVirtualMemory>
  <jsdl:UpperBoundedRange>30720.0</jsdl:Up-
    perBoundedRange>
</jsdl:IndividualVirtualMemory>
<jsdl:TotalVirtualMemory>
  <jsdl:UpperBoundedRange>245760.0</jsdl:Up-
    perBoundedRange>
</jsdl:TotalVirtualMemory>
</jsdl:Resources>

```

```
</jsdl:JobDescription>  
</jsdl:JobDefinition>
```

Example 10-6: File staging example:

```
qsub -N job6 -Wstagein=test.in:server@/home/test/test.in -Wstage-  
out=test.out:server@/home/test/test.out -- /bin/executable a b
```

JSDL document: If the `hpcbp_stage_protocol` attribute is set to `ftp:34` and the user with user-name `User1` has submitted the job, the HPCBP MOM generates the following JSDL document:

```
<?xml version="1.0" encoding="utf-8"?>
<jsdsl:JobDefinition xmlns:jsdsl="http://schemas.ggf.org/jsdsl/2005/11/
jsdl" xmlns:jsdsl-hpcpa="http://schemas.ggf.org/jsdsl/2006/07/jsdl-
hpcpa" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" >
  <jsdsl:JobDescription>
    <jsdsl:JobIdentification>
      <jsdsl:JobName>job6</jdsdl:JobName>
      <jsdsl:JobProject> PBS-TO-HPCBP v1</jdsdl:Job-
        Project>
    </jdsdl:JobIdentification>
    <jsdsl:Application>
      <jsdsl-hpcpa:HPCProfileApplication name="PBSto-
        HPCBP" >
        <jsdsl-hpcpa:Executable>/bin/execu-
          table</jdsdl-hpcpa:Executable>
        <jsdsl-hpcpa:Argument>a</jdsdl-hpcpa:Argu-
          ment>
        <jsdsl-hpcpa:Argument>b</jdsdl-hpcpa:Argu-
          ment>
      </jdsdl-hpcpa:HPCProfileApplication>
    </jdsdl:Application>
    <jsdsl:DataStaging>
      <jsdsl:FileName>test.in</jdsdl:FileName>
      <jsdsl:CreationFlag>overwrite</jdsdl:Creation-
        Flag>
      <jsdsl:Source>
        <jsdsl:URI>ftp://ser-
          ver:34/home/test/test.in</jdsdl:URI>
      </jdsdl:Source>
      <jsdsl:Credential>
```

```

    <jSDL:UsernameToken>
      <jSDL:Username>User1</jSDL:Username>
      <jSDL:password>User1</jSDL:password>
    </jSDL:UsernameToken>
  <jSDL:Credential>
</jSDL:DataStaging>
<jSDL:DataStaging>
  <jSDL:FileName>test.out</jSDL:FileName>
  <jSDL:CreationFlag>overwrite</jSDL:Creation-
    Flag>
  <jSDL:Target>
    <jSDL:URI>ftp://ser-
      ver:34/home/test/test.out</jSDL:URI>
  </jSDL:Target>
  <jSDL:Credential>
    <jSDL:UsernameToken>
      <jSDL:Username>User1</jSDL:Username>
      <jSDL:password>User1</jSDL:password>
    </jSDL:UsernameToken>
  <jSDL:Credential>
</jSDL:DataStaging>
</jSDL:JobDescription>
</jSDL:JobDefinition>

```

10.4 Caveats

10.4.1 Unsupported Commands

For a list of unsupported PBS commands, see section 7.7.2, "Unsupported Commands", on page 494 of the PBS Professional Administrator's Guide.

10.5 See Also

10.5.1 PBS Professional Manual Pages

See the following man pages:

- `qsub(1B)`
- `pbs_mom(1B)`

10.5.2 References

1. OGSA High Performance Computing Profile Working Group (OGSA-HPCP-WG) of the Open Grid Forum
<https://forge.gridforum.org/sf/projects/ogsa-hpcp-wg>
The HPC Basic Profile specification is GFD.114:
<http://www.ogf.org/documents/GFD.114.pdf>.
2. OGSA High Performance Computing Profile Working Group (OGSA-HPCP-WG) of the Open Grid Forum
<https://forge.gridforum.org/sf/projects/ogsa-hpcp-wg>
The HPC File Staging Profile Version 1.0:
<http://forge.ogf.org/sf/go/doc15024?nav=1>
3. OGSA Job Submission Description Language Working Group (JSDL - WG) of the Open Grid Forum
http://www.ogf.org/gf/group_info/view.php?group=jSDL-wg
The JSDL HPC Profile Application Extension, Version 1.0 is GFD 111:
<http://www.ogf.org/documents/GFD.111.pdf>
4. OGSA Usage Record Working Group (UR-WG) of the Open Grid Forum
The Usage Record - Format Recommendation is GFD.98
<http://www.ogf.org/documents/GFD.98.pdf>
5. Network Working Group, Uniform Resource Identifier (URI) : Generic Syntax
<http://www.rfc-editor.org/rfc/rfc3986.txt>

Appendix A: License Agreement

CAUTION!

PRIOR TO INSTALLATION OR USE OF THE SOFTWARE YOU MUST CONSENT TO THE FOLLOWING SOFTWARE LICENSE TERMS AND CONDITIONS BY CLICKING THE “I ACCEPT” BUTTON BELOW. YOUR ACCEPTANCE CREATES A BINDING LEGAL AGREEMENT BETWEEN YOU AND ALTAIR. IF YOU DO NOT HAVE THE AUTHORITY TO BIND YOUR ORGANIZATION TO THESE TERMS AND CONDITIONS, YOU MUST CLICK “I DO NOT ACCEPT” AND THEN HAVE AN AUTHORIZED PARTY IN THE ORGANIZATION THAT YOU REPRESENT ACCEPT THESE TERMS.

IF YOU, OR THE ORGANIZATION THAT YOU REPRESENT, HAS A MASTER SOFTWARE LICENSE AGREEMENT (“MASTER SLA”) ON FILE AT THE CORPORATE HEADQUARTERS OF ALTAIR ENGINEERING, INC. (“ALTAIR”), THE MASTER SLA TAKES PRECEDENCE OVER THESE TERMS AND SHALL GOVERN YOUR USE OF THE SOFTWARE.

MODIFICATION(S) OF THESE SOFTWARE LICENSE TERMS IS EXPRESSLY PROHIBITED. ANY ATTEMPTED MODIFICATION(S) WILL BE NONBINDING AND OF NO FORCE OR EFFECT UNLESS EXPRESSLY AGREED TO IN WRITING BY AN AUTHORIZED CORPORATE OFFICER OF ALTAIR. ANY DISPUTE RELATING TO THE VALIDITY OF AN ALLEGED MODIFICATION SHALL BE DETERMINED IN ALTAIR’S SOLE DISCRETION.

Altair Engineering, Inc. - Software License Agreement

THIS SOFTWARE LICENSE AGREEMENT, including any Additional Terms (together with the “Agreement”), shall be effective as of the date of YOUR acceptance of these software license terms and conditions (the “Effective Date”) and is between ALTAIR ENGINEERING, INC., 1820 E. Big Beaver Road, Troy, MI 48083-2031, USA, a Michigan corporation (“Altair”), and YOU, or the organization on whose behalf you have authority to accept these terms (the “Licensee”). Altair and Licensee, intending to be legally bound, hereby agree as follows:

1. DEFINITIONS. In addition to terms defined elsewhere in this Agreement, the following terms shall have the meanings defined below for purposes of this Agreement:

Additional Terms. Additional Terms are those terms and conditions which are determined by an Altair Subsidiary to meet local market conditions.

Documentation. Documentation provided by Altair or its resellers on any media for use with the Products.

Execute. To load Software into a computer's RAM or other primary memory for execution by the computer.

Global Zone: Software is licensed based on three Global Zones: the Americas, Europe and Asia-Pacific. When Licensee has Licensed Workstations located in multiple Global Zones, which are connected to a single License (Network) Server, a premium is applied to the standard Software License pricing for a single Global Zone.

ISV/Independent Software Vendor. A software company providing its products, (“ISV Software”) to Altair's Licensees through the Altair License Management System using Altair License Units.

License Log File. A computer file providing usage information on the Software as gathered by the Software.

License Management System. The license management system (LMS) that accompanies the Software and limits its use in accordance with this Agreement, and which includes a License Log File.

License (Network) Server. A network file server that Licensee owns or leases located on Licensee's premises and identified by machine serial number and/or HostID on the Order Form.

License Units. A parameter used by the LMS to determine usage of the Software permitted under this Agreement at any one time.

Licensed Workstations. Single-user computers located in the same Global Zone(s) that Licensee owns or leases that are connected to the License (Network) Server via local area network or Licensee's private wide-area network.

Maintenance Release. Any release of the Products made generally available by Altair to its Licensees with annual leases, or those with perpetual licenses who have an active maintenance agreement in effect, that corrects programming errors or makes other minor changes to the Software. The fees for maintenance and support services are included in the annual license fee but perpetual licenses require a separate fee.

Order Form. Altair's standard form in either hard copy or electronic format that contains the specific parameters (such as identifying Licensee's contracting office, License Fees, Software, Support, and License (Network) Servers) of the transaction governed by this Agreement.

Products. Products include Altair Software, ISV Software, and/or Suppliers' software; and Documentation related to all of the forgoing.

Proprietary Rights Notices. Patent, copyright, trademark or other proprietary rights notices applied to the Products, packaging or media.

Software. The Altair software identified in the Order Form and any Updates or Maintenance Releases.

Subsidiary. Subsidiary means any partnership, joint venture, corporation or other form of enterprise in which a party possesses, directly or indirectly, an ownership interest of fifty percent (50%) or greater, or managerial or operational control.

Suppliers. Any person, corporation or other legal entity which may provide software or documents which are included in the Software.

Support. The maintenance and support services provided by Altair pursuant to this Agreement.

Templates. Human readable ASCII files containing machine-interpretable commands for use with the Software.

Term. The term of licenses granted under this Agreement. Annual licenses shall have a 12-month term of use unless stated otherwise on the Order Form. Perpetual licenses shall have a term of twenty-five years. Maintenance agreements for perpetual licenses have a 12-month term.

Update. A new version of the Products made generally available by Altair to its Licensees that includes additional features or functionalities but is substantially the same computer code as the existing Products.

2. LICENSE GRANT. Subject to the terms and conditions set forth in this Agreement, Altair hereby grants Licensee, and Licensee hereby accepts, a limited, non-exclusive, non-transferable license to: a) install the Products on the License (Network) Server(s) identified on the Order Form for use only at the sites identified on the Order Form; b) execute the Products on Licensed Workstations in accordance with the LMS for use solely by Licensee's employees, or its onsite Contractors who have agreed to be bound by the terms of this Agreement, for Licensee's internal business use on Licensed Workstations within the Global Zone(s) as iden-

tified on the Order Form and for the term identified on the Order Form; c) make backup copies of the Products, provided that Altair's and its Suppliers' and ISV's Proprietary Rights Notices are reproduced on each such backup copy; d) freely modify and use Templates, and create interfaces to Licensee's proprietary software for internal use only using APIs provided that such modifications shall not be subject to Altair's warranties, indemnities, support or other Altair obligations under this Agreement; and e) copy and distribute Documentation inside Licensee's organization exclusively for use by Licensee's employees and its onsite Contractors who have agreed to be bound by the terms of this Agreement. A copy of the License Log File shall be made available to Altair automatically on no less than a monthly basis. In the event that Licensee uses a third party vendor for information technology (IT) support, the IT company shall be permitted to access the Software only upon its agreement to abide by the terms of this Agreement. Licensee shall indemnify, defend and hold harmless Altair for the actions of its IT vendor(s).

3. RESTRICTIONS ON USE. Notwithstanding the foregoing license grant, Licensee shall not do (or allow others to do) any of the following: a) install, use, copy, modify, merge, or transfer copies of the Products, except as expressly authorized in this Agreement; b) use any back-up copies of the Products for any purpose other than to replace the original copy provided by Altair in the event it is destroyed or damaged; c) disassemble, decompile or “unlock”, reverse translate, reverse engineer, or in any manner decode the Software or ISV Software for any reason; d) sublicense, sell, lend, assign, rent, distribute, publicly display or publicly perform the Products or Licensee's rights under this Agreement; e) allow use outside the Global Zone(s) or User Sites identified on the Order Form; f) allow third parties to access or use the Products such as through a service bureau, wide area network, Internet location or time-sharing arrangement except as expressly provided in Section 2(b); g) remove any Proprietary Rights Notices from the Products; h) disable or circumvent the LMS provided with the Products; or (i) link any software developed, tested or supported by Licensee or third parties to the Products (except for Licensee's own proprietary software solely for Licensee's internal use).

4. OWNERSHIP AND CONFIDENTIALITY. Licensee acknowledges that all applicable rights in patents, copyrights, trademarks, service marks, and trade secrets embodied in the Products are owned by Altair and/or its Suppliers or ISVs. Licensee further acknowledges that the Products, and all copies thereof, are and shall remain the sole and exclusive property of Altair and/or its Suppliers and ISVs. This Agreement is a license and not a sale of the Products. Altair retains all rights in the Products not expressly granted to Licensee herein. Licensee acknowledges that the Products are confidential and constitute valuable assets and trade secrets of Altair and/or its Suppliers and ISVs. Licensee agrees to take the same precautions necessary to protect and maintain the confidentiality of the Products as it does to protect its own information of a confidential nature but in any event, no less than a reasonable degree of care, and shall not disclose or make them available to any person or entity except as expressly provided in this Agreement. Licensee shall promptly notify Altair in the event any unauthorized person obtains access to the Products. If Licensee is required by any governmental

authority or court of law to disclose Altair's or its ISV's or its Suppliers' confidential information, then Licensee shall immediately notify Altair before making such disclosure so that Altair may seek a protective order or other appropriate relief. Licensee's obligations set forth in Section 3 and Section 4 of this Agreement shall survive termination of this Agreement for any reason. Altair's Suppliers and ISVs, as third party beneficiaries, shall be entitled to enforce the terms of this Agreement directly against Licensee as necessary to protect Supplier's intellectual property or other rights.

Altair and its resellers providing support and training to licensed end users of the Products shall keep confidential all Licensee information provided to Altair in order that Altair may provide Support and training to Licensee. Licensee information shall be used only for the purpose of assisting Licensee in its use of the licensed Products. Altair agrees to take the same precautions necessary to protect and maintain the confidentiality of the Licensee information as it does to protect its own information of a confidential nature but in any event, no less than a reasonable degree of care, and shall not disclose or make them available to any person or entity except as expressly provided in this Agreement.

5. MAINTENANCE AND SUPPORT. Maintenance. Altair will provide Licensee, at no additional charge for annual licenses and for a maintenance fee for paid-up licenses, with Maintenance Releases and Updates of the Products that are generally released by Altair during the term of the licenses granted under this Agreement, except that this shall not apply to any Term or Renewal Term for which full payment has not been received. Altair does not promise that there will be a certain number of Updates (or any Updates) during a particular year. If there is any question or dispute as to whether a particular release is a Maintenance Release, an Update or a new product, the categorization of the release as determined by Altair shall be final. Licensee agrees to install Maintenance Releases and Updates promptly after receipt from Altair. Maintenance Releases and Updates are subject to this Agreement. Altair shall only be obligated to provide support and maintenance for the most current release of the Software and the most recent prior release. **Support.** Altair will provide support via telephone and email to Licensee at the fees, if any, as listed on the Order Form. If Support has not been procured for any period of time for paid-up licenses, a reinstatement fee shall apply. Support consists of responses to questions from Licensee's personnel related to the use of the then-current and most recent prior release version of the Software. Licensee agrees to provide Altair with sufficient information to resolve technical issues as may be reasonably requested by Altair. Licensee agrees to the best of its abilities to read, comprehend and follow operating instructions and procedures as specified in, but not limited to, Altair's Documentation and other correspondence related to the Software, and to follow procedures and recommendations provided by Altair in an effort to correct problems. Licensee also agrees to notify Altair of a programming error, malfunction and other problems in accordance with Altair's then current problem reporting procedure. If Altair believes that a problem reported by Licensee may not be due to an error in the Software, Altair will so notify Licensee. Questions must be directed to Altair's specially designated telephone support numbers and email addresses. Support will also be available via email at Internet addresses designated by Altair. Support is available

Monday through Friday (excluding holidays) from 8:00 a.m. to 5:00 p.m. local time in the Global Zone where licensed, unless stated otherwise on the Order Form. **Exclusions.** Altair shall have no obligation to maintain or support (a) altered, damaged or Licensee-modified Software, or any portion of the Software incorporated with or into other software not provided by Altair; (b) any version of the Software other than the current version of the Software or the immediately prior release of the Software; (c) problems caused by Licensee's negligence, abuse or misapplication of Software other than as specified in the Documentation, or other causes beyond the reasonable control of Altair; or (d) Software installed on any hardware, operating system version or network environment that is not supported by Altair. Support also **excludes** configuration of hardware, non- Altair Software, and networking services; consulting services; general solution provider related services; and general computer system maintenance.

6. WARRANTY AND DISCLAIMER. Altair warrants for a period of ninety (90) days after Licensee initially receives the Software that the Software will perform under normal use substantially as described in then current Documentation. Supplier software included in the Software and ISV Software provided to Licensee shall be warranted as stated by the Supplier or the ISV. Copies of the Suppliers' and ISV's terms and conditions of warranty are available on the Altair Support website. Support services shall be provided in a workmanlike and professional manner, in accordance with the prevailing standard of care for consulting support engineers at the time and place the services are performed.

ALTAIR DOES NOT REPRESENT OR WARRANT THAT THE PRODUCTS WILL MEET LICENSEE'S REQUIREMENTS OR THAT THEIR OPERATION WILL BE UNINTERRUPTED OR ERROR-FREE, OR THAT IT WILL BE COMPATIBLE WITH ANY PARTICULAR HARDWARE OR SOFTWARE. ALTAIR EXCLUDES AND DISCLAIMS ALL EXPRESS AND IMPLIED WARRANTIES NOT STATED HEREIN, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT. THE ENTIRE RISK FOR THE PERFORMANCE, NON-PERFORMANCE OR RESULTS OBTAINED FROM USE OF THE PRODUCTS RESTS WITH LICENSEE AND NOT ALTAIR. ALTAIR MAKES NO WARRANTIES WITH RESPECT TO THE ACCURACY, COMPLETENESS, FUNCTIONALITY, SAFETY, PERFORMANCE, OR ANY OTHER ASPECT OF ANY DESIGN, PROTOTYPE OR FINAL PRODUCT DEVELOPED BY LICENSEE USING THE PRODUCTS.

7. INDEMNITY. Altair will defend and indemnify, at its expense, any claim made against Licensee based on an allegation that the Software infringes a patent or copyright ("Claim"); provided, however, that this indemnification does not include claims which are based on Supplier software or ISV software, and that Licensee has not materially breached the terms of this Agreement, Licensee notifies Altair in writing within ten (10) days after Licensee first learns of the Claim; and Licensee cooperates fully in the defense of the claim. Altair shall have sole control over such defense; provided, however, that it may not enter into

any settlement binding upon Licensee without Licensee's consent, which shall not be unreasonably withheld. If a Claim is made, Altair may modify the Software to avoid the alleged infringement, provided however, that such modifications do not materially diminish the Software's functionality. If such modifications are not commercially reasonable or technically possible, Altair may terminate this Agreement and refund to Licensee the prorated license fee that Licensee paid for the then current Term. Perpetual licenses shall be pro-rated over a 36-month term. Altair shall have no obligation under this Section 7, however, if the alleged infringement arises from Altair's compliance with specifications or instructions prescribed by Licensee, modification of the Software by Licensee, use of the Software in combination with other software not provided by Altair and which use is not specifically described in the Documentation, and if Licensee is not using the most current version of the Software, if such alleged infringement would not have occurred except for such exclusions listed here. This section 7 states Altair's entire liability to Licensee in the event a Claim is made. No indemnification is made for Supplier and/or ISV Software.

8. LIMITATION OF REMEDIES AND LIABILITY. Licensee's exclusive remedy (and Altair's sole liability) for Software that does not meet the warranty set forth in Section 6 shall be, at Altair's option, either (i) to correct the nonconforming Software within a reasonable time so that it conforms to the warranty; or (ii) to terminate this Agreement and refund to Licensee the license fees that Licensee has paid for the then current Term for the nonconforming Software; provided, however that Licensee notifies Altair of the problem in writing within the applicable Warranty Period when the problem first occurs. Any corrected Software shall be warranted in accordance with Section 6 for ninety (90) days after delivery to Licensee. The warranties hereunder are void if the Software has been improperly installed, misused, or if Licensee has violated the terms of this Agreement.

Altair's entire liability for all claims arising under or related in any way to this Agreement (regardless of legal theory), shall be limited to direct damages, and shall not exceed, in the aggregate for all claims, the license and maintenance fees paid under this Agreement by Licensee in the 12 months prior to the claim on a prorated basis, except for claims under Section 7. **ALTAIR AND ITS SUPPLIERS AND ISVS SHALL NOT BE LIABLE TO LICENSEE OR ANYONE ELSE FOR INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES ARISING HEREUNDER (INCLUDING LOSS OF PROFITS OR DATA, DEFECTS IN DESIGN OR PRODUCTS CREATED USING THE SOFTWARE, OR ANY INJURY OR DAMAGE RESULTING FROM SUCH DEFECTS, SUFFERED BY LICENSEE OR ANY THIRD PARTY) EVEN IF ALTAIR OR ITS SUPPLIERS OR ITS ISVS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.** Licensee acknowledges that it is solely responsible for the adequacy and accuracy of the input of data, including the output generated from such data, and agrees to defend, indemnify, and hold harmless Altair and its Suppliers and ISVs from any and all claims, including reasonable attorney's fees, resulting from, or in connection with Licensee's use of the Software. No

action, regardless of form, arising out of the transactions under this Agreement may be brought by either party against the other more than two (2) years after the cause of action has accrued, except for actions related to unpaid fees.

9. UNITED STATES GOVERNMENT RESTRICTED RIGHTS. This section applies to all acquisitions of the Products by or for the United States government. By accepting delivery of the Products except as provided below, the government or the party procuring the Products under government funding, hereby agrees that the Products qualify as “commercial” computer software as that term is used in the acquisition regulations applicable to this procurement and that the government's use and disclosure of the Products is controlled by the terms and conditions of this Agreement to the maximum extent possible. This Agreement supersedes any contrary terms or conditions in any statement of work, contract, or other document that are not required by statute or regulation. If any provision of this Agreement is unacceptable to the government, Vendor may be contacted at Altair Engineering, Inc., 1820 E. Big Beaver Road, Troy, MI 48083-2031; telephone (248) 614-2400. If any provision of this Agreement violates applicable federal law or does not meet the government's actual, minimum needs, the government agrees to return the Products for a full refund.

For procurements governed by DFARS Part 227.72 (OCT 1998), the Software, except as described below, is provided with only those rights specified in this Agreement in accordance with the Rights in Commercial Computer Software or Commercial Computer Software Documentation policy at DFARS 227.7202-3(a) (OCT 1998). For procurements other than for the Department of Defense, use, reproduction, or disclosure of the Software is subject to the restrictions set forth in this Agreement and in the Commercial Computer Software - Restricted Rights FAR clause 52.227-19 (June 1987) and any restrictions in successor regulations thereto.

Portions of Altair's PBS Professional Software and Documentation are provided with RESTRICTED RIGHTS. Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subdivision(c)(1)(ii) of the rights in the Technical Data and Computer Software clause in DFARS 252.227-7013, or in subdivision (c)(1) and (2) of the Commercial Computer Software-Restricted Rights clause at 48 CFR52.227-19, as applicable.

10. CHOICE OF LAW AND VENUE. This Agreement shall be governed by and construed under the laws of the state of Michigan, without regard to that state's conflict of laws principles except if the state of Michigan adopts the Uniform Computer Information Transactions Act drafted by the National Conference of Commissioners of Uniform State Laws as revised or amended as of June 30, 2002 (“UCITA”) which is specifically excluded. This Agreement shall not be governed by the United Nations Convention on Contracts for the International Sale of Goods, the application of which is expressly excluded. Each Party waives its right to a jury trial in the event of any dispute arising under or relating to this Agreement. Each party agrees that money damages may not be an adequate remedy for breach

of the provisions of this Agreement, and in the event of such breach, the aggrieved party shall be entitled to seek specific performance and/or injunctive relief (without posting a bond or other security) in order to enforce or prevent any violation of this Agreement.

11. [RESERVED]

12. Notice. All notices given by one party to the other under the Agreement or these Additional Terms shall be sent by certified mail, return receipt requested, or by overnight courier, to the respective addresses set forth in this Agreement or to such other address either party has specified in writing to the other. All notices shall be deemed given upon actual receipt.

Written notice shall be made to:

Altair: **Licensee Name & Address:**
Altair Engineering, Inc. _____
1820 E. Big Beaver Rd _____
Troy, MI 48083 _____
Attn: Tom M. Perring Attn: _____

13. TERM. For annual licenses, or Support provided for perpetual licenses, renewal shall be automatic for each successive year (“Renewal Term”), upon mutual written execution of a new Order Form. All charges and fees for each Renewal Term shall be set forth in the Order Form executed for each Renewal Term. All Software licenses procured by Licensee may be made coterminous at the written request of Licensee and the consent of Altair.

14. TERMINATION. Either party may terminate this Agreement upon thirty (30) days prior written notice upon the occurrence of a default or material breach by the other party of its obligations under this Agreement (except for a breach by Altair of the warranty set forth in Section 8 for which a remedy is provided under Section 10; or a breach by Licensee of Section 5 or Section 6 for which no cure period is provided and Altair may terminate this Agreement immediately) if such default or breach continues for more than thirty (30) days after receipt of such notice. Upon termination of this Agreement, Licensee must cease using the Software and, at Altair's option, return all copies to Altair, or certify it has destroyed all such copies of the Software and Documentation.

15. GENERAL PROVISIONS. Export Controls. Licensee acknowledges that the Products may be subject to the export control laws and regulations of the United States and other countries, and any amendments thereof. Licensee agrees that Licensee will not directly or indirectly export the Products into any country or use the Products in any manner except in compliance with all applicable U.S. and other countries export laws and regulations. **Notice.** All notices given by one party to the other under this Agreement shall be sent by certified mail, return receipt requested, or by overnight courier, to the respective addresses set forth in

this Agreement or to such other address either party has specified in writing to the other. All notices shall be deemed given upon actual receipt. **Assignment.** Neither party shall assign this Agreement without the prior written consent of other party, which shall not be unreasonably withheld. All terms and conditions of this Agreement shall be binding upon and inure to the benefit of the parties hereto and their respective successors and permitted assigns. **Waiver.** The failure of a party to enforce at any time any of the provisions of this Agreement shall not be construed to be a waiver of the right of the party thereafter to enforce any such provisions. **Severability.** If any provision of this Agreement is found void and unenforceable, such provision shall be interpreted so as to best accomplish the intent of the parties within the limits of applicable law, and all remaining provisions shall continue to be valid and enforceable. **Headings.** The section headings contained in this Agreement are for convenience only and shall not be of any effect in constructing the meanings of the Sections. **Modification.** No change or modification of this Agreement will be valid unless it is in writing and is signed by a duly authorized representative of each party. **Conflict.** In the event of any conflict between the terms of this Agreement and any terms and conditions on a Licensee Purchase Order or comparable document, the terms of this Agreement shall prevail. Moreover, each party agrees any additional terms on any Purchase Order or comparable document other than the transaction items of (a) item(s) ordered; (b) pricing; (c) quantity; (d) delivery instructions and (e) invoicing directions, are not binding on the parties. In the event of a conflict between the terms of this Agreement, and the Additional Terms, the Agreement shall take precedence. **Entire Agreement.** This Agreement, the Additional Terms, and the Order Form(s) attached hereto constitute the entire understanding between the parties related to the subject matter hereto, and supersedes all proposals or prior agreements, whether written or oral, and all other communications between the parties with respect to such subject matter. This Agreement may be executed in one or more counterparts, all of which together shall constitute one and the same instrument. **Execution.** Copies of this Agreement executed via original signatures, facsimile or email shall be deemed binding on the parties.

Index

A

activereq 94
addreq 94
allreq 94

B

batch 11

C

client commands 5
closerm 94
commands 5
configrm 94
Cray
 SV1 vii
 T3e vii
credential 29

D

downrm 94

E

Executor 5

F

flushreq 94

fullresp 94

G

getreq 94
Globus 5
 MOM 5

J

Job
 Executor (MOM) 5
 Scheduler 5
 Server 4

M

manager 5
 commands 5
MOM 5
monitoring 3

N

NASA vii
NQS 3

O

openrm 94
operator

Index

- commands 5
- P**
- pbs_alterjob 31
- pbs_asyruntimejob 52
- PBS_BATCH_AsyrunJob 12
- PBS_BATCH_AuthenUser 11
- PBS_BATCH_Commit 11
- PBS_BATCH_Connect 11
- PBS_BATCH_CopyFiles 11
- PBS_BATCH_CopyFiles_Cred 11
- PBS_BATCH_DeleteJob 11
- PBS_BATCH_DeleteResv 12
- PBS_BATCH_DelFiles 11
- PBS_BATCH_DelFiles_Cred 11
- PBS_BATCH_Disconnect 11
- PBS_BATCH_FailOver 11
- PBS_BATCH_GSS_Context 11
- PBS_BATCH_HoldJob 11
- PBS_BATCH_JobCred 11
- PBS_BATCH_JobObit 11
- PBS_BATCH_jobscript 11
- PBS_BATCH_LocateJob 11
- PBS_BATCH_Manager 11
- PBS_BATCH_MessJob 11
- PBS_BATCH_ModifyJob 11
- PBS_BATCH_MoveJob 11
- PBS_BATCH_MvJobFile 11
- PBS_BATCH_OrderJob 11
- PBS_BATCH_QueueJob 11
- PBS_BATCH_RdytoCommit 11
- PBS_BATCH_RegistDep 11
- PBS_BATCH_ReleaseJob 11
- PBS_BATCH_ReleaseResc 11
- PBS_BATCH_Rerun 11
- PBS_BATCH_Rescq 11
- PBS_BATCH_ReserveResc 11
- PBS_BATCH_RunJob 11
- PBS_BATCH_SelectJobs 11
- PBS_BATCH_SelStat 11
- PBS_BATCH_Shutdown 11
- PBS_BATCH_SignalJob 11
- PBS_BATCH_StageIn 11
- PBS_BATCH_StatusJob 11
- PBS_BATCH_StatusNode 11
- PBS_BATCH_StatusQue 11
- PBS_BATCH_StatusResv 12
- PBS_BATCH_StatusSvr 11
- PBS_BATCH_SubmitResv 11
- PBS_BATCH_TrackJob 12
- pbs_connect 33
- pbs_default 35
- pbs_deljob 36
- pbs_delresv 37
- pbs_disconnect 38
- pbs_geterrmsg 39
- pbs_holdjob 40
- pbs_locjob 41
- pbs_manager 42
- pbs_module(7B) 112
- pbs_mom 5
- pbs_movejob 45
- pbs_msgjob 46
- pbs_orderjob 47
- pbs_rerunjob 48
- pbs_rescrelease 49
- pbs_resreserve 49
- pbs_rlsjob 51
- pbs_runjob 52
- pbs_sched 5
- pbs_selectjob 54
- pbs_selstat 56
- pbs_server 4
- pbs_sigjob 59
- pbs_stagein 60
- pbs_stathook(3B) 121
- pbs_stathost 63
- pbs_statjob 61
- pbs_statnode 63
- pbs_statque 65
- pbs_statresv 67
- pbs_statsched 69
- pbs_statsserver 71
- pbs_statvnode 63

Index

pbs_submit 73
pbs_submitsrv 75
pbs_tclapi 98
pbs_tclsh 97
pbs_terminate 77
pbs_wish 97
POSIX 5

Q

queuing 3
Quick Start Guide ix

R

rpp_bind 80
rpp_close 80
rpp_eom 80
rpp_flush 80
rpp_getaddr 80
rpp_getc 80
rpp_io 80
rpp_open 79, 80
rpp_poll 80
rpp_putc 80
rpp_rcommit 80
rpp_read 80
rpp_shutdown 80
rpp_terminate 80
rpp_wcommit 80
rpp_write 80

S

Scheduler 5
scheduling 3
Server 4
system daemons 4

T

TCL 97
tm_atnode 86
tm_attach 86
tm_finalize 86

tm_init 86
tm_kill 86
tm_nodeinfo 86
tm_notify 86
tm_obit 86
tm_poll 86
tm_publish 86
tm_rescinfo 86
tm_spawn 86
tm_subscribe 86
tm_taskinfo 86

U

user
 commands 4, 5
User Guide ix

W

workload management 3

Index
